

Dynamic Intransitive Noninterference Revisited

Sebastian Eggert
 Institut für Informatik,
 Kiel University,
 Kiel, Germany.
 e-mail: research@sebastian-eggert.de

Ron van der Meyden
 School of Computer Science and Engineering,
 UNSW Australia,
 Sydney NSW 2052, Australia.
 e-mail: meyden@cse.unsw.edu.au

January 21, 2016

Abstract

The paper studies dynamic information flow security policies in an automaton-based model. Two semantic interpretations of such policies are developed, both of which generalize the notion of TA-security [van der Meyden ESORICS 2007] for static intransitive noninterference policies. One of the interpretations focuses on information flows permitted by policy edges, the other focuses on prohibitions implied by absence of policy edges. In general, the two interpretations differ, but necessary and sufficient conditions are identified for the two interpretations to be equivalent. Sound and complete proof techniques are developed for both interpretations. Two applications of the theory are presented. The first is a general result showing that access control mechanisms are able to enforce a dynamic information flow policy. The second is a simple capability system motivated by the Flume operating system.

1 Introduction

Covert channels, i.e., avenues for information flow unintended by the designers of a system, present a significant risk for systems that are required to enforce confidentiality, since they constitute vulnerabilities that an attacker may exploit in order to exfiltrate secrets. This has motivated a body of research that aims to provide means by which systems can be assured to be free of covert channels. At the highest levels of assurance, one aims for formal verification based on rigorous models of the system. This requires formal definitions of what it means for a

system to comply with a policy that specifies the intended information flows, and prohibits the unintended information flows.

How to give a formal semantics to such policies that is suitable for formal verification has been a topic of study since the 1970's [10]. *Noninterference policies* [15] are a form of policy syntax in the form of a graph over a set of security levels, or domains, with an edge representing that information is permitted to flow from the source of the edge to the destination. Conversely, absence of an edge can be understood as prohibiting (direct) information flow. In the original formulation of such policies, corresponding to classical multi-level security policies based on a lattice of security levels, the graph was transitive, but there has more recently been a body of work on *intransitive noninterference* policies, which do not require transitivity. Here, a domain is best understood as a component within a security architecture, with edges describing permitted direct flows of information, and paths corresponding to permitted indirect flows of information. In particular, this understanding underlies MILS security [2, 36], an approach to high assurance systems development in which security architectures comprised of trusted and untrusted components are enforced using a variety of mechanisms including separation kernels. At its most abstract level, the architecture can be understood as an intransitive noninterference policy. Several large-scale efforts [6, 19, 24, 27, 31] have formally verified that intransitive noninterference policies hold in a number of separation kernels.

The proper semantics of intransitive noninterference policies has proved to be a subtle matter. Rushby [30] improved an initial proposal by Haigh and Young [17], giving semantics to policies by means of an *intransitive purge* function. Van der Meyden [34] has recently argued that there is a subtle weakness in this intransitive purge based semantics. He provides an example that demonstrates that this semantics may classify as secure a system that contains flows of information, based on the ordering of events, that seem to contradict an intuitive understanding of the policy. He proposes an alternative, stronger, definition of security, *TA-security*, that correctly classifies this example as insecure. Moreover, TA-security is shown to provide a better fit for Rushby's proof theory. Rushby shows that this proof theory is sound for his definition, but fails to provide a completeness result. Van der Meyden explains this, by showing that it is both sound and complete for the stronger notion of TA-security. This completeness result is obtained both for Rushby's unwinding proof method and for a variant of Rushby's results demonstrating that access control systems provide a sound method for enforcement of intransitive noninterference policies.

This line of work has concerned static policies, in which the permitted flows of information do not change over time. Such policies are too restrictive for many practical purposes: the permitted flows of information in a system could change as a result of policy changes, system reconfiguration, in response to detection of vulnerabilities or attacks in progress, as a result of personnel changes, promotions, revocations of privileges, delegations, transfers of capabilities, and for many other reasons.

A full understanding of the formal underpinnings of security architecture should accommodate such dynamic policy changes, and it is therefore of inter-

est to have a satisfactory theory of dynamic intransitive noninterference policies. To date, only a few works have studied this question: Leslie [22] and Eggert et. al. [11]. Both are based on generalizations of Rushby’s “intransitive purge” operator and, in the case of static policies, are equivalent to Rushby’s semantics. Consequently, these works are subject to the weakness of this semantics identified by van der Meyden. In particular, they misclassify van der Meyden’s example. Leslie provides an unwinding proof theory for her definition that is equivalent to Rushby’s unwinding in the static case, hence is incomplete. Eggert et. al. do provide a complete proof theory and a complexity characterization for their definition.

Our main contribution in this paper is to develop alternate semantics for dynamic intransitive noninterference policies that generalize the more adequate notion of TA-security in the static case. Indeed, we show (in Section 3) that there are several candidates for such a semantics since, in the dynamic setting, it turns out that the information flows permitted by policy edges are not always the complement of the information flows prohibited by the absence of policy edges. This leads to two distinct semantics, one of which prioritizes the permissions and the other of which prioritizes the prohibitions. A further complexity to policy in the dynamic setting is the question of what knowledge the agents in the system are permitted to have of the state of the policy itself. However, this consideration is also helpful: we show that *locality*, a natural condition on agent knowledge of the policy, provides necessary and sufficient conditions for the permissive and the prohibitive interpretations of the policy to be equivalent. Reference to agent knowledge also enables us to characterize precisely the maximal information permitted to be known to each agent in the system at moment of time, in terms of a concrete representation of the information that may have been transmitted to the agent. We show by some examples that our definitions improve on those of Leslie [22] and Eggert et. al. [11], not just in the treatment of static policies (van der Meyden’s example justifying TA-security), but also in the classification of some simple dynamic policies.

We then (in Section 4) develop proof methods that may be used to show that a system is secure, with a focus on local policies, so that the technique applies to both of our definitions of security. We first generalize the classical *unwinding* proof method [16] to our new definitions. This is used (in Section 5) to generalize, to a dynamic setting, Rushby’s [30] results on enforcement of a static intransitive noninterference policy by means of access control settings. In both cases, we demonstrate not just soundness but also completeness of the proof method.

As an application of the theory, we establish (in Section 6) security for a capability system, motivated by the Flume operating system [21]. We show that the capability system enforces a naturally associated dynamic intransitive noninterference policy. This result gives an illustration of how the general theory can be applied to a concrete system.

2 Basic Definitions

In this section we give the basic definitions concerning the systems model about which we will make security judgements, and introduce the syntax of dynamic intransitive noninterference policies. The following section develops several semantics for these policies in this systems model.

2.1 Systems

We model systems as deterministic automata in which the actions of multiple agents cause state transitions. Agents obtain information about the state of system, and each other's actions, by making observations of the state of the system.

A *signature* is a tuple $\langle D, \text{dom}, A \rangle$, where D is a finite set of *domains*, (which we may also refer to as *agents*), A is a finite set of *actions*, and $\text{dom}: A \rightarrow D$, assigns each action to a domain (agent). Intuitively, $\text{dom}(a)$ is the domain within which the action a is performed (the agent performing that action). We write A^* for the set of finite sequences of elements of the set A and denote the empty sequence by ϵ .

An *automaton* for a signature $\langle D, \text{dom}, A \rangle$ is a tuple $\langle S, s_0, \rightarrow \rangle$, where S is a set of states (not necessarily finite), $s_0 \in S$ is the *initial state*, and $\rightarrow \subseteq S \times A \times S$ is a labelled transition relation. This relation is required to be *input-enabled*, in the sense that for all states $s \in S$ and actions $a \in A$, there exists $t \in S$ such that $(s, a, t) \in \rightarrow$. We also assume that the relation is deterministic in the sense that if $(s, a, t) \in \rightarrow$ and $(s, a, t') \in \rightarrow$ then $t = t'$. Under these assumptions, we may write $s \cdot a$ for the unique t such that $(s, a, t) \in \rightarrow$.

We make several uses of the general notion of automaton. One is the machine model that we analyze for security. Formally, a *system* for a signature $\langle D, \text{dom}, A \rangle$ is an automaton $\langle S, s_0, \rightarrow \rangle$ for the signature, together with an observation function $\text{obs}: D \times S \rightarrow O$. Intuitively, $\text{obs}(u, s)$ is the observation made in domain u (by agent u) when the system is in state s . Since the domain u will typically be fixed in a context of application, but the state s will vary, we write $\text{obs}_u(s)$ for $\text{obs}(u, s)$. We write systems in the form $M = \langle S, s_0, \rightarrow, D, \text{dom}, A, \text{obs} \rangle$ when we wish to make all the relevant components explicit.

The assumption that the transition relation is input-enabled has often been made in the literature. It prevents enablement of actions being a source of information flow, and intuitively represents that it is always possible for an agent to attempt to perform an action (even if it would have no actual effect). A situation s where an agent is not able to successfully perform an action can be represented in an input-enabled model by means of a transition (s, a, s) that leaves the state unchanged. If it is desired to model that the agent is able to observe that the action is not enabled, this can be encoded in the observation that the agent makes in either the source or destination state of the transition. (This leaves the modelling of the way that enablement of actions becomes known to agents in the hands of the system modeller, as opposed to many process

algebraic semantics, e.g., those using bisimulation, that would mandate that enablement of actions is always observable.)

The restriction to deterministic systems is made in part because the theory of intransitive noninterference is better understood in deterministic systems [30, 34] than in nondeterministic systems [13]. However, in many cases, nondeterminism can be represented in deterministic systems by using additional agents to model the source of nondeterminism. (This does require that such additional agents be included in the policy as well as in the systems model, but that can be argued to add clarity to the meaning of the policy.)

A *run* of a system is a sequence $s_0 a_1 s_1 \dots a_n s_n$, where $n \geq 0$ and the $s_i \in S$ are states (with s_0 the initial state of the system) and the $a_i \in A$ are actions, such that for $i = 1 \dots n$ we have $s_i = s_{i-1} \cdot a_i$. We call the sequence $\alpha = a_1 \dots a_n \in A^*$ the *trace* of the run. Conversely, since the system is input-enabled and deterministic, every sequence $\alpha \in A^*$ is the trace of a unique run. We may therefore refer to runs by their corresponding traces. For a sequence $\alpha = a_1 \dots a_n \in A^*$ and a state $s \in S$, we write $s \cdot \alpha$ for the state reached after executing the sequence of actions α starting at s , i.e., the state t_n , obtained inductively by taking $t_0 = s$ and $t_i = t_{i-1} \cdot a_i$, for $i = 1 \dots n$.

Two systems $M = \langle S, s_o, \rightarrow, \text{obs} \rangle$ and $M' = \langle S', s'_o, \rightarrow', \text{obs}' \rangle$ for the same signature $\langle D, \text{dom}, A \rangle$ are *bisimilar* if for all $u \in D$ and $\alpha \in A^*$, we have $\text{obs}_u(s_o \cdot \alpha) = \text{obs}'_{u'}(s'_o \cdot \alpha)$. (We remark that we do not need the usual relational definition of bisimilarity because our systems are deterministic.) It is easily checked that all the definitions of security that follow are preserved under bisimilarity.

The *unfolding* of a system M for signature $\langle D, \text{dom}, A \rangle$ is defined to be the system $\text{unfold}(M)$ for the same signature, with states A^* , initial state ϵ , transition relation such that $\alpha \cdot a = \alpha a$ for all $\alpha \in A^*$ and $a \in A$, and observation functions $\text{obs}_u(\alpha) = \text{obs}_u^M(s_o \cdot \alpha)$. The system $\text{unfold}(M)$ is easily seen to be bisimilar to M , so $\text{unfold}(M)$ can be thought of as an equivalent implementation of the system M .

2.2 Static Policies

A *static policy* describes the permitted and prohibited flows of information between domains. Formally, a static policy for the signature $\langle D, \text{dom}, A \rangle$ is a reflexive binary relation \rightharpoonup (possibly intransitive) on the set of domains D . Intuitively, $u \rightharpoonup v$ means that information is permitted to flow from domain u to domain v , and conversely, $u \not\rightharpoonup v$ means that information flows from domain u to domain v are prohibited.

Policies are assumed to be reflexive because, intuitively, nothing can be done to prevent information flowing from a domain to itself, so this is allowed by default. Since \rightharpoonup can be taken to be the edge set of a directed graph with vertex set D , we may use terminology from graph theory when we describe properties of a policy. For example, we may call $u \rightharpoonup v$ an “edge” of the policy \rightharpoonup , and talk about “paths” in the relation \rightharpoonup .

We remark that the system model introduced above appears to lack a notion of internal (silent) transitions, but these can be handled by adding a systems

domain Sys , such that each internal transition corresponds to some action of domain Sys . This representation also requires adding an edge $Sys \rightarrow u$ to the policy for all other domains u .

A significant body of work already exists in relation to static policies, e.g., [15, 30, 34, 29]. Since our focus in this paper is to find semantics for dynamic policies that generalizes the semantics of van der Meyden [34] for static intransitive noninterference policies, we reiterate here just that semantics in order to motivate what follows, and refer the reader to the literature for more detailed discussion. The semantics involves for each agent u a function \mathbf{ta}_u from the set of traces, such that for a trace α , the value $\mathbf{ta}_u(\alpha)$ intuitively represents a concrete encoding of the maximal amount of information that the agent is permitted to have about α . The function \mathbf{ta}_u transforms a trace α into a binary tree, which contains less information about the order of the actions than the trace does. This function is defined inductively by¹ $\mathbf{ta}_u(\epsilon) = \epsilon$ and for every $\alpha \in A^*$ and $a \in A$ by

$$\mathbf{ta}_u(\alpha a) = \begin{cases} (\mathbf{ta}_u(\alpha), \mathbf{ta}_{\text{dom}(a)}(\alpha), a) & \text{if } \text{dom}(a) \rightarrow u \\ \mathbf{ta}_u(\alpha) & \text{otherwise} \end{cases}.$$

Intuitively, in case $\text{dom}(a) \not\rightarrow u$, information from $\text{dom}(a)$ is not permitted to flow to u , so the occurrence of the action a causes no change in the maximal information $\mathbf{ta}_u(\alpha)$ that may be possessed by u after α . However, if $\text{dom}(a) \rightarrow u$, then the occurrence of a after α causes the maximal information that agent u is permitted to have to increase, by adding to its prior information $\mathbf{ta}_u(\alpha)$ the maximal information $\mathbf{ta}_{\text{dom}(a)}(\alpha)$ that the agent performing a is permitted to have after α , as well as the fact that the action a has occurred. Using the functions \mathbf{ta}_u , we may now declare a system to be secure, in the sense of complying with the information flow policy \rightarrow , as follows:

Definition 1 *A system M is TA-secure with respect to policy \rightarrow , if for all agents u and $\alpha, \beta \in A^*$, if $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\beta)$ then $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \beta)$.*

Intuitively, this says that for all agents u , the observation $\text{obs}_u(s_0 \cdot \alpha)$ contains no more information than the maximal permitted information $\mathbf{ta}_u(\alpha)$.

In general, when considering the deductive capabilities of agents in computer security, it is appropriate to assume that an adversary will reason based not just on their current observation, but also using their past observations. This suggests that we should consider a *perfect recall* attacker. In the case of TA-security, a definition stated using a perfect recall attacker is equivalent to one stated using just the final observation of traces [34]. We formulate this result in a more general way, so that it also applies to the other definitions we state in this paper.

¹ Note that the functions \mathbf{ta}_u depend on the policy \rightarrow , but to avoid clutter, this is suppressed in the notation, here and for similar definitions later in the paper. We may add the policy to the notation in contexts containing several policies, e.g., by writing $\mathbf{ta}_u^{\rightarrow}(\alpha)$ when the policy used is \rightarrow .

First, to capture a perfect recall attacker, we define the *view* of an agent u inductively, by $\mathbf{view}_u(\epsilon) = \mathbf{obs}_u(s_0)$ and

$$\mathbf{view}_u(\alpha a) = \begin{cases} \mathbf{view}_u(\alpha) a \mathbf{obs}_u(s_0 \cdot \alpha a) & \text{if } \mathbf{dom}(a) = u \\ \mathbf{view}_u(\alpha) \circ \mathbf{obs}_u(s_0 \cdot \alpha a) & \text{otherwise} \end{cases}$$

where every $\alpha \in A^*$ and $a \in A$. Here, $\sigma \circ x$ denotes the absorptive concatenation of an element x of a set X to a string whose final element is in X , defined by $\epsilon x = x$, and $\sigma y \circ x = \sigma y x$ if $x \neq y$ and $\sigma y \circ x = \sigma y$ if $y = x$. The reason we apply this type of concatenation is to capture the asynchronous behaviour of the system, where agents do not necessarily have access to a global clock. Intuitively, agent u is aware of its own actions, so these are always available in its perfect recall view. This is captured by the first clause of the definition of \mathbf{view}_u . However, when another agent performs an action, this may or may not cause a change in u 's observation. If it does, then u detects the change, and the new observation is added to its view. However, if there is no change of observation, then there is no change to the view, since the agent is assumed to operate asynchronously, so that it does not know for what length of time it makes each observation.

Consider an indexed collection of functions $f = \{f_u\}_{u \in D}$, each with domain A^* , and defined by $f_u(\epsilon) = \epsilon$ and

$$f_u(\alpha a) = \begin{cases} (f_u(\alpha), g(\alpha, a, u)) & \text{if } C(\alpha, a, u) \\ f_u(\alpha) & \text{otherwise} \end{cases}$$

for some boolean condition C and function g , where $\alpha \in A^*$, $a \in A$ and $u \in D$. Note that the collection \mathbf{ta} is in this pattern, with $g(\alpha, a, u) = (f_{\mathbf{dom}(a)}(\alpha), a)$ and $C(\alpha, a, u) = \mathbf{dom}(a) \rightarrow u$. We say that f is *self-aware* if $\mathbf{dom}(a) = u$ implies $C(\alpha, a, u)$, and $g(\alpha, a, u) = g(\beta, b, u)$ implies $a = b$. Intuitively, in the self-aware case, $f_u(\alpha)$ encodes at least a record of all the actions of domain u that occur in α .

Say that a system is *f-secure* if for all $\alpha, \beta \in A^*$ with $f_u(\alpha) = f_u(\beta)$, we have $\mathbf{obs}_u(s_0 \cdot \alpha) = \mathbf{obs}_u(s_0 \cdot \beta)$. An alternative definition with the \mathbf{view} function instead of \mathbf{obs} is: a system is *f-view-secure* if for all $\alpha, \beta \in A^*$ with $f_u(\alpha) = f_u(\beta)$, we have $\mathbf{view}_u(\alpha) = \mathbf{view}_u(\beta)$. These definitions are equivalent, subject to self-awareness of f :

Lemma 1 *Suppose that f is self-aware. Then a system M is f -secure iff it is f -view-secure.*

Proof: Assume first that M is not f -view-secure. Then there are $u \in D$ and $\alpha, \beta \in A^*$ of minimal combined length with $f_u(\alpha) = f_u(\beta)$ and $\mathbf{view}_u(\alpha) \neq \mathbf{view}_u(\beta)$. At least one of α and β is not the empty trace, suppose that it is α and let $\alpha = \alpha' a$ for some $\alpha' \in A^*$ and $a \in A$. Then there are two cases:

- *Case 1:* $f_u(\alpha' a) = f_u(\alpha')$. Then $f_u(\alpha') = f_u(\beta)$, so from the minimality of α and β it follows that $\mathbf{view}_u(\alpha') = \mathbf{view}_u(\beta)$ and hence $\mathbf{view}_u(\alpha') \neq$

$\text{view}_u(\alpha'a)$. Since $C(\alpha, a, u)$ is false, we have by self-awareness that $\text{dom}(a) \neq u$, so $\text{view}_u(\alpha'a) = \text{view}_u(\alpha')\text{obs}_u(s_0 \cdot \alpha'a)$. Therefore, we have $\text{obs}_u(s_0 \cdot \alpha'a) \neq \text{obs}_u(s_0 \cdot \beta)$.

- *Case 2:* $f_u(\alpha'a) \neq f_u(\alpha')$. We can assume that $\beta = \beta'b$ for some $\beta' \in A^*$ and $b \in A$ with $f_u(\beta'b) \neq f_u(\beta')$ since otherwise, we proceed with the first case with the roles of α and β swapped. From the inductive definition of f_u and self-awareness, it follows that $a = b$ and $f_u(\alpha') = f_u(\beta')$. By the minimality of α and β , we have $\text{view}_u(\alpha') = \text{view}_u(\beta')$. There are two cases of the definitions of $\text{view}_u(\alpha'a)$ and $\text{view}_u(\beta'a)$, depending on whether $\text{dom}(a) = u$, but in either case, it follows that $\text{obs}_u(s_0 \cdot \alpha'a) \neq \text{obs}_u(s_0 \cdot \beta'b)$, since the only difference in these sequences can be in the final observation.

The other direction of the proof follows directly from the definition of the **view** function. \square

Since the definition of $f = \mathbf{ta}$ is easily seen to be self-aware (using reflexivity of the policy), it follows from Lemma 1 that TA-security and **ta-view**-security are equivalent.

2.3 Dynamic Policies

Our main concern in this paper is with *dynamic intransitive information flow policies*, which generalize static policies by allowing the edges of the policy to depend on the actions that have been performed in the system. Formally, a *dynamic policy* is a relation $\mapsto \subseteq D \times A^* \times D$, such that $(u, \alpha, u) \in \mapsto$ for all $\alpha \in A^*$ and all $u \in D$. We write $\alpha \models u \mapsto v$ when $(u, \alpha, v) \in \mapsto$. Intuitively, $\alpha \models u \mapsto v$ says that, after the sequence of actions α have been performed, information may flow from domain u to domain v . Thus, $\mapsto_\alpha = \{(u, v) \mid (u, \alpha, v) \in \mapsto\}$ is the (static) policy that applies after the actions α have been performed.

Dynamic policies can be represented using automata. We say that a tuple $\langle P, \mapsto' \rangle$ consisting of an automaton $P = \langle S, s_0, \rightarrow \rangle$ and a relation $\mapsto' \subseteq D \times S \times D$ represents a dynamic policy \mapsto , if for all $\alpha \in A^*$, domains $u, v \in D$, and state $s \in S$ with $s = s_0 \cdot \alpha$, we have $\alpha \models u \mapsto v$ iff $(u, s, v) \in \mapsto'$. (In this case we also write $s \models u \mapsto' v$.) Every dynamic policy has such a representation (with an infinite number of states), since we may take $S = A^*$ and $\rightarrow = \{(\alpha, a, \alpha a) \mid \alpha \in A^*, a \in A\}$. We may call the policy *finite state* if it has a representation with S finite.

Let $\langle D, \text{dom}, A \rangle$ be a signature. A *policy enhanced system* for this signature is an automaton $\mathcal{A} = \langle S, s_0, \rightarrow \rangle$ that is equipped with a relation $\mapsto \subseteq D \times S \times D$ and an observation function **obs** with domain $D \times S$. Given a system M and a dynamic policy \mapsto' for the signature, we say that the policy enhanced system $\langle \mathcal{A}, \mapsto, \text{obs} \rangle$ *encodes* M and \mapsto' if $\langle \mathcal{A}, \text{obs} \rangle$ is bisimilar to M and $\langle \mathcal{A}, \mapsto \rangle$ represents \mapsto' .

For every pair consisting of a system M and a dynamic policy \mapsto for the same signature, we can construct a policy enhanced system that encodes M and \mapsto .

Given an automaton representation $\langle \mathcal{A}_P, \rightarrow' \rangle$ of \rightarrow where $\mathcal{A}_P = \langle S^P, s_0^P, \rightarrow^P \rangle$, and a system $M = \langle \mathcal{A}_M, \text{obs}^M \rangle$ where $\mathcal{A}_M = \langle S^M, s_0^M, \rightarrow^M \rangle$ for the same signature, we define the product automaton $\mathcal{A}_M \times \mathcal{A}_P = \langle S, s_0, \rightarrow \rangle$, where $S = S^M \times S^P$, $s_0 = (s_0^M, s_0^P)$ and $((s, p), a, (s', p')) \in \rightarrow$ iff $(s, a, s') \in \rightarrow^M$ and $(p, a, p') \in \rightarrow^P$. The automaton $\mathcal{A}_M \times \mathcal{A}_P$ may be equipped with an observation function obs with domain S and a policy relation $\rightarrow'' \subseteq D \times S \times D$ by defining $\text{obs}(u, (s, p)) = \text{obs}^M(u, s)$ and $\rightarrow'' = \{(u, (s, p), v) \mid u \rightarrow_p v\}$. It is then straightforward to show that the policy enhanced system $\langle \mathcal{A}_M \times \mathcal{A}_P, \text{obs}, \rightarrow'' \rangle$ encodes M and \rightarrow .

It will often be convenient, in presenting examples, to use a diagram of a policy enhanced system in order to present the policy and the system together as a single automaton. Figure 1 gives an example of this presentation. We note the following conventions to be applied in interpreting these diagrams. The systems represented are input-enabled, but we elide self loops corresponding to edges of the form (s, a, s) in order to reduce clutter. States are partitioned into three components. The top component gives the name of the state. The middle component depicts the static policy applicable at that state. The bottom component gives partial information about the observations that the agents make at the state: typically, only one agent's information is relevant to the discussion, and we elide the observations of the other agents in the system.

There is a natural order on policies. Consider the binary relation \leq on policies with respect to a given signature $\langle D, \text{dom}, A \rangle$, defined by $\rightarrow \leq \rightarrow'$ iff for all $\alpha \in A^*$ and $u, v \in D$, we have $\alpha \models u \rightarrow v$ implies $\alpha \models u \rightarrow' v$. It is easily seen that this relation partially orders the policies with respect to the signature. Intuitively, if $\rightarrow \leq \rightarrow'$ then \rightarrow places more restrictions on the flow of information in a system than does \rightarrow' . This intuition can be supported for both a permissive and a prohibitive reading of policies. On a permissive reading, $\rightarrow \leq \rightarrow'$ intuitively says that every situation in which a flow of information is explicitly permitted by \rightarrow is one where the flow of information is explicitly permitted by \rightarrow' . Thus, \rightarrow is more restrictive than \rightarrow' in the sense of having fewer explicitly permitted flows of information. On a prohibitive reading, $\rightarrow \leq \rightarrow'$ says (contrapositively) that every situation where \rightarrow' explicitly prohibits a flow of information is one where \rightarrow also explicitly prohibits that flow of information. Thus, \rightarrow is more restrictive than \rightarrow' in the sense of having more explicitly prohibited flows of information. We may therefore gloss $\rightarrow \leq \rightarrow'$ as stating that \rightarrow is *more restrictive* than \rightarrow' .

2.4 Logic of Knowledge

It will be convenient to formulate some of our definitions using formulas from a logic of knowledge. Given a signature $\langle D, \text{dom}, A \rangle$, we work with formulas ϕ expressed using the following grammar:

$$\phi ::= u \rightarrow v \mid \phi \wedge \phi \mid \neg \phi \mid K_u \phi \mid D_G \phi$$

where $u, v \in D$ are domains and $G \subseteq D$ is a set of domains. Intuitively, the atomic propositions of the logic are assertions of the form $u \rightarrow v$ concerning

the static policy holding at a particular point of time. We write $Prop$ for the set of atomic propositions. The logic contains the usual boolean operators for conjunction and negation, and we freely use other boolean operators that can be defined using these, e.g., we write $\phi \Rightarrow \psi$ for $\neg(\phi \wedge \neg\psi)$. The formula $K_u\phi$ intuitively says that domain u *knows* that ϕ holds, and $D_G\phi$ says that it is *distributed knowledge* to the group G that ϕ holds, i.e., the group G would be able to deduce ϕ if they were to pool all the information held by the members of the group.

The semantics of the logic is a standard Kripke semantics for epistemic logic [14]: formulas are interpreted in Kripke structures $\mathcal{K} = \langle W, \{\sim_u\}_{u \in D}, \models \rangle$ where W is a set, for each $u \in D$, we have an equivalence relation \sim_u on W , and $\models \subseteq W \times Prop$ is a binary relation.

Intuitively, W is a set of worlds, representing possible situations in a system of interest. The equivalence relation \sim_u intuitively corresponds to indistinguishability of worlds to an agent: $w \sim_u w'$ will hold when agent u has the same information available to it when it is in situation w as it has in situation w' . For a world w and an atomic proposition p , the relation $w \models p$ represents that the proposition p is true at the world w . We may also write this as $\mathcal{K}, w \models p$ to make the structure \mathcal{K} explicit. In turn, this relation can be extended to a satisfaction relation $\mathcal{K}, w \models \phi$ for arbitrary formulas, by means of the following recursion:

$$\begin{aligned} \mathcal{K}, w \models \phi_1 \wedge \phi_2 & \text{ if } \mathcal{K}, w \models \phi_1 \text{ and } \mathcal{K}, w \models \phi_2 \\ \mathcal{K}, w \models \neg\phi & \text{ if not } \mathcal{K}, w \models \phi \\ \mathcal{K}, w \models K_u\phi & \text{ if } \mathcal{K}, w' \models \phi \text{ for all } w' \in W \text{ with } w \sim_u w' \\ \mathcal{K}, w \models D_G\phi & \text{ if } \mathcal{K}, w' \models \phi \text{ for all } w' \in W \text{ with } w \sim_u w' \\ & \text{ for all } u \in G \end{aligned}$$

In our applications, we will work with the set of worlds $W = A^*$, i.e., worlds will be traces for the signature of interest. The basic relation $w \models p$ will be the relation $\alpha \models u \rightarrow v$ from some dynamic policy for this signature. We discuss the equivalence relations we use later.

3 Semantics for dynamic policies

In the setting of dynamic policies, several subtle issues arise that a suitable definition of security needs to take into account.

One is that policies can be interpreted with a focus on positive or negative edges. One can read $\alpha \models u \rightarrow v$ as stating a permission: in state α , actions of u may pass information from u to v (even if u or v does not know that $\alpha \models u \rightarrow v$). Nothing about the policy or system may override this permission, and an agent cannot be sanctioned for having caused the information flow. Alternately, and more restrictively, one could focus on the converse: and treat $\alpha \models u \not\rightarrow v$ as stating a prohibition: in state α , information may not flow from u to v . The interaction between various such prohibitions may have the effect that even where there is an edge $\alpha \models u \rightarrow v$, it is in fact prohibited for information to




Figure 1: Conflicting permissive and prohibitive interpretations

flow from u to v because of derived prohibitions. The following example suggests that these two readings of the policy may be in conflict.

Example 1: Consider the system in Figure 1. Intuitively, domain P is a policy authority that controls the policy between domains A and B by means of its action p , and a is an action of domain A .

This system displays a conflict in the transition from state s_1 to state s_2 , between the permissive reading of the edge $A \mapsto B$ and the prohibitive reading of the edge $P \not\mapsto B$. On the one hand, according to a permissive reading of the policy, the fact that the action a causes a change in the observation in domain B is not an insecurity, because the policy explicitly permits a flow of information from domain A to B in state s_1 , and what B learns in state s_2 is that A has performed action a . On the other hand, in no state of the policy is there an edge from domain P to domain B . Even the union of all possible policy states does not contain a path from P to B . Thus, we expect from a prohibitive reading of the policy that there should be no flow of information from domain P to domain B . On this reading, the transition from s_1 to s_2 displays an insecurity in the system, since on making observation 1 domain B can deduce not just that action a has occurred, but also that it has been preceded by action p in domain P . Thus, the system contains a prohibited flow of information from P to B , and should be judged to be insecure. \square

There are several attitudes one can take in the face of such apparent conflicts within a policy. One is to allow explicit permissions to take precedence over prohibitions. Another is to require that prohibitions take precedence over permissions. Finally, we can restrict policies to instances that do not contain such conflicts. We will consider all these possibilities.

For static policies, it is safe to assume that the policy is known to all agents in the system: indeed, it can be assumed to be common knowledge. A second subtlety for dynamic policies is that the current policy state depends on past actions, and not all agents will be permitted to know the entire past. This implies that there may be situations where some agents do not have complete information about the current policy state. This causes difficulty for any agent that is required to enforce compliance with the policy. For example, in the policy of Example 1, agent A is never permitted to know whether P has performed action p , so cannot be relied upon to enforce a policy restriction on the transfer

of information about its actions to agent B , as applies in the initial policy state, but not in state s_1 . We will approach this issue below by interpreting the policy in a way takes into account what agents know about the policy state.

3.1 Permissive Interpretation

We first formulate a definition that captures a permissive reading of a dynamic policy, while generalizing TA-security. The methodology underlying the notion of TA-security is to construct a concrete representation of the maximal information that each agent is permitted to have after each sequence of actions, and then to state that a system is secure if it has no more than that permitted information. For dynamic policies, with a permissive interpretation of edges $\alpha \models u \mapsto v$, this leads to the following inductive definition of an operator \mathbf{ta}^\diamond . It is defined for every domain u , trace $\alpha \in A^*$ and action $a \in A$ by $\mathbf{ta}_u^\diamond(\epsilon) = \epsilon$ and

$$\mathbf{ta}_u^\diamond(\alpha a) = \begin{cases} (\mathbf{ta}_u^\diamond(\alpha), \mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha), a) & \text{if } \alpha \models \text{dom}(a) \mapsto u \\ \mathbf{ta}_u^\diamond(\alpha) & \text{otherwise} \end{cases}.$$

We also obtain a derived equivalence relation \sim_u^\diamond on A^* for each $u \in D$, defined by $\alpha \sim_u^\diamond \beta$ iff $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$.

Intuitively, $\alpha \models \text{dom}(a) \mapsto u$ states that it is permitted, after trace α , for $\text{dom}(a)$ to send information to u , and $\text{dom}(a)$ exercises this by sending to u all the information that it has (viz. $\mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha)$), as well as the fact that it is performing the action a . The agent u receives this information and appends it to the information that it already has (viz., $\mathbf{ta}_u^\diamond(\alpha)$.) In case $\alpha \models \text{dom}(a) \not\mapsto u$, on the other hand, there is no transfer of information from $\text{dom}(a)$ to u , since the policy does not permit this.

If we take \mathbf{ta}^\diamond as a formalization of the maximal information permitted to an agent, we get the following definition of security, in the same pattern as TA-security.

Definition 2 *A system M is \mathbf{ta}^\diamond -secure with respect to a policy \mapsto if for all domains u and sequences of actions α, β , if $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$ then $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \beta)$.*

We remark that the definition of \mathbf{ta}^\diamond is self-aware (since by reflexivity $\text{dom}(a) = u$ implies $\alpha \models \text{dom}(a) \mapsto u$ for all α) so it follows from Lemma 1 that \mathbf{ta}^\diamond -security and \mathbf{ta}^\diamond -view-security are equivalent.

The following example shows that \mathbf{ta}^\diamond -security takes a permissive interpretation of Example 1.

Example 2: Treating the diagram of Figure 1 as a definition of the policy, we obtain that for any sequence α of the form a^* or a^*p , we have $\mathbf{ta}_B^\diamond(\alpha) = \epsilon$. On the other hand $\mathbf{ta}_B^\diamond(pa) = (\epsilon, \epsilon, a)$, so every sequence β such that $\mathbf{ta}_B^\diamond(pa) = \mathbf{ta}_B^\diamond(\beta)$ has the property that it contains a p followed later by an a . Thus, if we interpret

$\mathbf{ta}_B^\diamond(pa)$ as a concrete representation of what B is allowed to know after the sequence pa , we conclude that B is permitted to know that there has been an occurrence of action p . Indeed, for any sequence β with $\mathbf{ta}_B^\diamond(pa) = \mathbf{ta}_B^\diamond(\beta)$ we have $s_0 \cdot \beta = s_2 = s_0 \cdot pa$, so there is not a witness to insecurity of the system at the sequence pa . The reader may confirm that this system is in fact \mathbf{ta}^\diamond -secure.

We remark here that we have $\mathbf{ta}_A^\diamond(p) = \epsilon = \mathbf{ta}_A^\diamond(\epsilon)$ and $\mathbf{ta}_B^\diamond(p) = \epsilon = \mathbf{ta}_B^\diamond(\epsilon)$, but $\mathbf{ta}_B^\diamond(pa) = (\epsilon, \epsilon, a) \neq \epsilon = \mathbf{ta}_B^\diamond(a)$. That is, in general, the value $\mathbf{ta}_B^\diamond(\alpha a)$ depends on more than the values $\mathbf{ta}_B^\diamond(\alpha)$ and $\mathbf{ta}_A^\diamond(\alpha)$ and the action a ; additional information (namely, does $\alpha \models \text{dom}(a) \rightarrow u$ hold) about the sequence α is required to determine $\mathbf{ta}_B^\diamond(\alpha a)$. We discuss the significance of this below. \square

3.2 Prohibitive Interpretation

Example 2 shows that if we would like instead to prioritize prohibitions over permissions when interpreting a policy, we need a different representation of the maximal information permitted to an agent than that provided by the function \mathbf{ta}^\diamond . It is not completely obvious what should be the deduced prohibitions that prevent the use of an edge $\alpha \models \text{dom}(a) \rightarrow u$ as a justification for the transmission of information from $\text{dom}(a)$ to u in the transition from state α to state αa .

However, we may apply some general constraints to suggest a definition. First of all, note that for static policies, the function \mathbf{ta} has the property that the information available to domain u immediately after an action a performed by domain $\text{dom}(a)$, should depend only on the action a and the information possessed by domains u and $\text{dom}(a)$. The function \mathbf{ta}^\diamond does not have this property: see the remark at the end of Example 2.

A general way to ensure that the state of information of domain u after action a depends only on the prior states of information of domains u and $\text{dom}(a)$, is to condition the definition of the maximal information permitted to be known to u on the distributed knowledge of the group $\{u, \text{dom}(a)\}$. This suggests the following inductive definition of a function \mathbf{ta}^\square , similar to \mathbf{ta} , defined by $\mathbf{ta}_u^\square(\epsilon) = \epsilon$ and if $\alpha \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \rightarrow u)$, then

$$\mathbf{ta}_u^\square(\alpha a) = (\mathbf{ta}_u^\square(\alpha), \mathbf{ta}_{\text{dom}(a)}^\square(\alpha), a)$$

and else $\mathbf{ta}_u^\square(\alpha a) = \mathbf{ta}_u^\square(\alpha)$.

The formula $D_{\{\text{dom}(a), u\}}(\text{dom}(a) \rightarrow u) \Rightarrow \text{dom}(a) \rightarrow u$ is a validity of epistemic logic, so this definition strengthens the condition under which information is transmitted in the definition of \mathbf{ta}^\diamond . Intuitively, this definition permits transmission of $\text{dom}(a)$'s information to u only when it is distributed knowledge to u and $\text{dom}(a)$ that $\text{dom}(a) \rightarrow u$, i.e., there is not the possibility, so far as these domains jointly know, that there is an edge $\text{dom}(a) \not\rightarrow u$ that prohibits the transmission of information from $\text{dom}(a)$ to u . Thus, the definition takes the point of view that transmission is permitted whenever the agents would jointly be able to confirm that there is not an explicit prohibition to the transmission.

As it stands, the definition of \mathbf{ta}^\square is incomplete, because in order to interpret the condition $\alpha \models D_{\{\mathbf{dom}(a), u\}}(\mathbf{dom}(a) \rightarrow u)$, we need an appropriate Kripke structure in order to evaluate the distributed knowledge operator. We may take A^* to be the set of worlds of this structure, and derive satisfaction of atomic propositions from the policy, but it remains to specify equivalence relations \sim_v for $v \in D$. Intuitively, an equivalence $\alpha \sim_v \beta$ holds when agent v is not able to distinguish between the traces α and β : whenever the system is in state α , agent v considers it possible that the system is in state β . What these equivalence relations should be is not immediately apparent. We take the approach of identifying some constraints on these equivalence relations.

First, consider the general constraint, that the way agent u 's state of information is updated when action a is performed, should depend only on the prior state of information of u and $\mathbf{dom}(a)$. This can be captured at the level of the relations \sim_v by means of the following constraint:

(WSC) For all traces α, β , and actions a , if $\alpha \sim_u \beta$ and $\alpha \sim_{\mathbf{dom}(a)} \beta$ then $\alpha a \sim_u \beta a$.

(We call this constraint WSC since it is essentially the same as Rushby's condition of *Weak Step Consistency* [30].) Although it does not mention the policy explicitly, this constraint is not inconsistent with the state of the policy being a factor in how u 's knowledge is updated. However, it requires that the way that the policy is taken into account depends only on information jointly available to domains u and $\mathbf{dom}(a)$.

Next, to factor in the prohibitions to information transfer implied by the policy, note that one case where the policy clearly prohibits transfer of information from a domain $\mathbf{dom}(a)$ to u when action a is performed after sequence α is where $\alpha \models \mathbf{dom}(a) \not\rightarrow u$. Thus, domain u 's state of information in state αa should be the same as its information in state α in this case. This suggests the following constraint:

(DLR) For all traces α , and actions a , if $\alpha \models \mathbf{dom}(a) \not\rightarrow u$ then $\alpha a \sim_u \alpha$.

The nomenclature DLR arises from the fact that this is essentially a dynamic generalization of Rushby's condition of *Locally Respects* [30]. A slightly stronger alternative of DLR is the following:

(DLR') For all traces α , and actions a , if there exists a trace β such that $\alpha \sim_u \beta$ and $\alpha \sim_{\mathbf{dom}(a)} \beta$ and $\beta \models \mathbf{dom}(a) \not\rightarrow u$ then $\alpha a \sim_u \alpha$.

In fact, in the context of WSC, this is equivalent to DLR.

Proposition 1 *If WSC then DLR iff DLR'.*

Proof: We assume WSC and DLR and show DLR'. (The other direction is straightforward.) Suppose that $\alpha \sim_u \beta$ and $\alpha \sim_{\mathbf{dom}(a)} \beta$ and $\beta \models \mathbf{dom}(a) \not\rightarrow u$. We need to show $\alpha a \sim_u \alpha$. By WSC, we have $\alpha a \sim_u \beta a$. Since $\beta \models \mathbf{dom}(a) \not\rightarrow u$, we have by DLR that $\beta a \sim_u \beta$. Two applications of transitivity now yield that $\alpha a \sim_u \alpha$, as required. \square

Suppose that we now accept WSC and DLR as reasonable constraints on the equivalence relations \sim_u corresponding to the maximal knowledge that an agent u is permitted to have. Indexed collections of equivalence relations may be partially ordered by pointwise containment. (That is, we take $\{\sim_u\}_{u \in D} \leq \{\sim'_u\}_{u \in D}$ if $\sim_u \subseteq \sim'_u$ for all $u \in D$.) Given the dynamic policy \mapsto , let $\{\sim_u^{unw}\}_{u \in D}$ be the smallest collection of equivalence relations that is consistent with WSC and DLR. Intuitively, these smallest relations are the ones that allow for the *maximal* amount of information flow in the system that is consistent with the constraints. (The smaller the equivalence class, the larger the agent's amount of knowledge.) We call these relations the *unwinding relations*. The following fact will be useful below.

Lemma 2 *Suppose $\mapsto \leq \mapsto'$ and let $\{\sim_u^{unw}\}_{u \in D}$ and $\{\approx_u^{unw}\}_{u \in D}$ be the unwinding relations for the policies \mapsto and \mapsto' , respectively. Then $\approx_u^{unw} \subseteq \sim_u^{unw}$ for all $u \in D$.*

Proof: Since $\text{dom}(a) \not\mapsto' u$ implies $\text{dom}(a) \not\mapsto u$, there are fewer cases of basic facts $\alpha a \sim_u^{unw} \alpha$ than basic facts $\alpha a \approx_u^{unw} \alpha$ derivable using DLR. By monotonicity of WSC and the rules for an equivalence relation, it follows that $\approx_u^{unw} \subseteq \sim_u^{unw}$ for all $u \in D$. \square

We now take the relations $\{\sim_u^{unw}\}_{u \in D}$ as the equivalence relations in the Kripke structure needed to interpret the distributed knowledge operator in the definition of the functions \mathbf{ta}_u^\square . We call the derived security notion associated to these functions \mathbf{ta}^\square -security. (Note that since, by reflexivity, if $\text{dom}(a) = u$ then $D_{\{u, \text{dom}(a)\}}(\text{dom}(a) \mapsto u)$ is valid in this Kripke structure, these functions are self-aware, so again, by Lemma 1, \mathbf{ta}^\square -security and \mathbf{ta}^\square -view-security are equivalent.) For $u \in D$ we also write \sim_u^\square for the relation on traces defined by $\alpha \sim_u^\square \beta$ iff $\mathbf{ta}_u^\square(\alpha) = \mathbf{ta}_u^\square(\beta)$.

However, we note that we could also take the unwinding relations directly as the basis for another definition of security:

Definition 3 *A system is unwinding-secure with respect to a policy \mapsto if, for all traces α, β and domain u , we have $\alpha \sim_u^{unw} \beta$ implies $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \beta)$.*

In fact, it turns out that this does not give a new definition of security. The following result shows that the unwinding relations correspond precisely to the information flow modelled by the function \mathbf{ta}^\square .

Theorem 3.1 *We have $\sim_u^{unw} = \sim_u^\square$ for all $u \in D$.*

Proof: We establish the containments $\sim_u^{unw} \subseteq \sim_u^\square$ and $\sim_u^\square \subseteq \sim_u^{unw}$.

For $\sim_u^{unw} \subseteq \sim_u^\square$, we argue by induction on the number of steps of a derivation of the statement $\alpha \sim_u^{unw} \beta$ using rules WSC, DLR and the rules for an equivalence relation. The cases for reflexivity, transitivity and symmetry are direct from the fact that \sim_u^\square has these properties. For the DLR case, suppose that $\alpha \models \text{dom}(a) \not\mapsto u$, so that we have $\alpha a \sim_u^{unw} \alpha$. By validity of the formula $D_G(\phi) \Rightarrow \phi$, we have also $\alpha \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \not\mapsto u)$, so by the definition of

\mathbf{ta}^\square we have $\alpha a \sim_u^\square \alpha$, as required. For the WSC case, suppose that $\alpha \sim_u^{unw} \beta$ and $\alpha \sim_{\text{dom}(a)}^{unw} \beta$, so that $\alpha a \sim_u^{unw} \beta a$ by WSC. By the inductive hypothesis, we have $\alpha \sim_u^\square \beta$ and $\alpha \sim_{\text{dom}(a)}^\square \beta$, i.e., $\mathbf{ta}_u^\square(\alpha) = \mathbf{ta}_u^\square(\beta)$ and $\mathbf{ta}_{\text{dom}(a)}^\square(\alpha) = \mathbf{ta}_{\text{dom}(a)}^\square(\beta)$. Moreover, it also follows that $\alpha \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \not\rightarrow u)$ iff $\beta \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \not\rightarrow u)$, so that the same case of the definition of \mathbf{ta}^\square is selected in both $\mathbf{ta}_u^\square(\alpha a)$ and $\mathbf{ta}_u^\square(\beta a)$. It follows that $\mathbf{ta}_u^\square(\alpha a) = \mathbf{ta}_u^\square(\beta a)$, i.e., $\alpha a \sim_{\text{dom}(a)}^\square \beta a$, as required.

For $\sim_u^\square \subseteq \sim_u^{unw}$, we argue by induction on the combined length of α and β that $\alpha \sim_{\text{dom}(a)}^\square \beta$ implies $\alpha \sim_{\text{dom}(a)}^{unw} \beta$. The base case of $\alpha = \beta = \epsilon$ is trivial by reflexivity. By symmetry, it suffices to consider the case of $\alpha = \alpha a$ and β , where the result has already been established for sequences of shorter combined length. We suppose that $\alpha \sim_{\text{dom}(a)}^\square \beta$, i.e., $\mathbf{ta}_u^\square(\alpha a) = \mathbf{ta}_u^\square(\beta)$, and show that $\alpha \sim_{\text{dom}(a)}^{unw} \beta$. We consider two cases, corresponding to the cases of the definition of \mathbf{ta}^\square .

In the first case, we have $\alpha \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \not\rightarrow u)$, so that $\mathbf{ta}_u^\square(\alpha) = \mathbf{ta}_u^\square(\alpha a)$, and hence $\mathbf{ta}_u^\square(\alpha) = \mathbf{ta}_u^\square(\beta)$. By induction, we have $\alpha \sim_u^{unw} \beta$. Since $\alpha \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \not\rightarrow u)$, there exists a trace γ such that $\alpha \sim_u^{unw} \gamma$ and $\alpha \sim_{\text{dom}(a)}^{unw} \gamma$ and $\gamma \models \text{dom}(a) \not\rightarrow u$. The relation \sim_u^{unw} satisfies WSC and DLR by definition, so by Proposition 1 it also satisfies DLR'. Thus, we obtain that $\alpha a \sim_u^{unw} \alpha$. It now follows by transitivity that $\alpha a \sim_u^{unw} \beta$, as required.

In the second case, we have $\alpha \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \rightarrow u)$, so that $\mathbf{ta}_u^\square(\alpha a) = (\mathbf{ta}_u^\square(\alpha), \mathbf{ta}_{\text{dom}(a)}^\square(\alpha), a)$. Assuming that $\mathbf{ta}_u^\square(\alpha a) = \mathbf{ta}_u^\square(\beta)$, it follows that $\beta \neq \epsilon$, so we may write $\beta = \beta' b$. If $\beta' \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \not\rightarrow u)$, then we may switch the roles of αa and $\beta' b$ and argue as above, so it suffices to consider the case where $\beta' \models D_{\{\text{dom}(a), u\}}(\text{dom}(a) \rightarrow u)$. In this case, we have $\mathbf{ta}_u^\square(\beta' b) = (\mathbf{ta}_u^\square(\beta'), \mathbf{ta}_{\text{dom}(b)}^\square(\beta'), b)$, so it follows from $\mathbf{ta}_u^\square(\alpha a) = \mathbf{ta}_u^\square(\beta)$ that $a = b$ and $\mathbf{ta}_u^\square(\alpha) = \mathbf{ta}_u^\square(\beta')$ and $\mathbf{ta}_{\text{dom}(a)}^\square(\alpha) = \mathbf{ta}_{\text{dom}(a)}^\square(\beta')$. That is, we have $\alpha \sim_u^{unw} \beta'$ and $\alpha \sim_{\text{dom}(a)}^{unw} \beta'$. By WSC, it follows that $\alpha a \sim_u^{unw} \beta' a = \beta$, as required. \square

It is immediate from Theorem 3.1 that unwinding security and \mathbf{ta}^\square -security are equivalent. Thus, these definitions lend support to each other and help to explain each other: the functions \mathbf{ta}^\square give an intuitive, concrete description of the information flows that are permitted when security is interpreted using the unwinding relations.

Moreover, Theorem 3.1 shows that \mathbf{ta}^\square -security provides a feasible general framework for security definitions. To instantiate this framework, we must provide a collection of equivalence relations $\{\sim_u\}_{u \in D}$ to be used to interpret the distributed knowledge operator in the definition. It is reasonable to ask that however we do so, it should be the case that the resulting relations $\{\sim_u^\square\}_{u \in D}$ are exactly the relations $\{\sim_u\}_{u \in D}$, so that the definition of \mathbf{ta}^\square is self-consistent.² Theorem 3.1 shows that this condition is in fact satisfiable.

There remains some circularity in this justification of the definitions of un-

²It is worth noting that the circularity here is similar to the circularity in the semantics of *knowledge-based programs* [14].

winding security and \mathbf{ta}^\square -security. In particular, these definitions are both based on the assumption that the constraints WSC and DLR are the only constraints that should be applied to the flow of information in order to satisfy the policy, and that the resulting relations give an acceptable notion of permitted agent knowledge. However, if a case can be made that further constraints should be placed on the semantics, then a comparison similar to that made above can be considered.

The next lemma shows that \mathbf{ta}^\square -security implies \mathbf{ta}^\diamond -security.

Lemma 3 *We have $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$ implies $\alpha \sim_u^{unw} \beta$ for all $\alpha, \beta \in A^*$ and all $u \in D$. Moreover, \mathbf{ta}^\square -security implies \mathbf{ta}^\diamond -security.*

Proof: By induction on the combined length of α and β . The base case is trivial by reflexivity. Consider the case of αa and β , and suppose $\mathbf{ta}_u^\diamond(\alpha a) = \mathbf{ta}_u^\diamond(\beta)$.

In case $\alpha \models \text{dom}(a) \not\rightarrow u$, we have $\mathbf{ta}_u^\diamond(\alpha a) = \mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$, so by induction we obtain $\alpha \sim_u^{unw} \beta$. Also, by DLR, we have $\alpha a \sim_u^{unw} \alpha$. It follows that $\alpha a \sim_u^{unw} \beta$, as required.

Alternately, if $\alpha \models \text{dom}(a) \rightarrow u$, then we have $\mathbf{ta}_u^\diamond(\alpha a) = (\mathbf{ta}_u^\diamond(\alpha), \mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha), a)$. It follows that $\beta \neq \epsilon$, so let $\beta = \beta' b$ for $b \in A$. In the case that $\beta' \models \text{dom}(b) \not\rightarrow u$, we may swap the roles of αa and $\beta' b$ and argue as in the previous paragraph. In case $\beta' \models \text{dom}(b) \rightarrow u$, we have $\mathbf{ta}_u^\diamond(\beta' b) = (\mathbf{ta}_u^\diamond(\beta'), \mathbf{ta}_{\text{dom}(b)}^\diamond(\beta'), b)$, and it follows that $a = b$ and $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta')$ and $\mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha) = \mathbf{ta}_{\text{dom}(a)}^\diamond(\beta')$. By induction, we obtain that $\alpha \sim_u^{unw} \beta'$ and $\alpha \sim_{\text{dom}(a)}^{unw} \beta'$. It now follows by WSC that $\alpha a \sim_u^{unw} \beta' a = \beta$, as required. \square

The following example shows that \mathbf{ta}^\square -security differs from \mathbf{ta}^\diamond -security.

Example 3: As already seen in Example 2, the system in Figure 1 is \mathbf{ta}^\diamond -secure. But, we have that $p \sim_B^{unw} \epsilon$ and $p \sim_A^{unw} \epsilon$ by WSC. By applying DLR, we obtain $pa \sim_B^{unw} a$. Using Theorem 3.1, we have $\mathbf{ta}_B^\square(pa) = \mathbf{ta}_B^\square(a)$, and since $\text{obs}_B(s_0 \cdot pa) \neq \text{obs}_B(s_0 \cdot a)$, the system is not \mathbf{ta}^\square -secure. \square

3.3 A sufficient condition for equivalence

The outcome of the discussion so far is that we have identified two definitions of security of dynamic policies, one of which (\mathbf{ta}^\diamond -security) takes a permissive interpretation of policies, and the other (\mathbf{ta}^\square -security, or equivalently, unwinding security) takes a prohibitive interpretation. As noted above, these two interpretations may be in conflict for some policies. It is therefore of interest to understand when a policy is free of this conflict. The following notion is useful in this regard.

Definition 4 *A policy \rightarrow is local with respect to a collection of equivalence relations $\{\approx_u\}_{u \in D}$ on A^* if, for all domains $u, v \in D$ and traces $\alpha, \beta \in A^*$, if $\alpha \approx_u \beta$ and $\alpha \approx_v \beta$ then $\alpha \models u \rightarrow v$ iff $\beta \models u \rightarrow v$.*

An equivalent statement is that

$$\alpha \models D_{\{u,v\}}(u \rightarrow v) \vee D_{\{u,v\}}(u \not\rightarrow v)$$

for all traces α and domains u, v , when the equivalence relation used to interpret knowledge is \approx -equivalence. That is, a pair of agents always have distributed knowledge of whether one may interfere with the other.

Example 4: The policy of Example 1 is not a local policy with respect to the relations $\{\sim_u^\diamond\}_{u \in D}$. For, we have $\epsilon \sim_A^\diamond p$ and $\epsilon \sim_B^\diamond p$, but $\epsilon \models A \not\rightarrow B$ and $p \models A \rightarrow B$. We present some examples of classes of local policies in Section 3.4. \square

We could instantiate the definition of locality with respect to either the equivalence relations $\{\sim_u^\diamond\}_{u \in D}$ or $\{\sim_u^\square\}_{u \in D}$ (equivalently, the unwinding relations $\{\sim_u^{unw}\}_{u \in D}$). In fact, these definitions prove to be equivalent. We first remark that, in general, the relations \sim_u^{unw} differ from the relations \sim_u^\diamond . In particular, $\alpha \sim_u^{unw} \beta$ does not imply $\alpha \sim_u^\diamond \beta$.

Example 5: In this example we proceed with the system of Figure 1. By DLR, we have from $\epsilon \models P \not\rightarrow B$ that $p \sim_B^{unw} \epsilon$ and from $\epsilon \models P \not\rightarrow A$ that $p \sim_A^{unw} \epsilon$. By WSC, we have $pa \sim_B^{unw} a$. However, the \mathbf{ta}^\diamond values of B differ: $\mathbf{ta}_B^\diamond(pa) = (\epsilon, \epsilon, a) \neq \epsilon = \mathbf{ta}_B^\diamond(a)$. \square

However, the relations \sim_u^{unw} and \sim_u^\diamond are identical for policies that are local with respect to the equivalence relations $\{\sim_u^\diamond\}_{u \in D}$.

Lemma 4 *Let \rightarrow be a local policy with respect to $\{\sim_u^\diamond\}_{u \in D}$. Then $\{\sim_u^\diamond\}_{u \in D} = \{\sim_u^{unw}\}_{u \in D}$.*

Proof: Suppose that the policy \rightarrow is local with respect to $\{\sim_u^\diamond\}_{u \in D}$. Then these relations satisfy WSC. For, suppose that $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$ and $\mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha) = \mathbf{ta}_{\text{dom}(a)}^\diamond(\beta)$. Note that, by locality, we have that $\alpha \models \text{dom}(a) \rightarrow v$ iff $\beta \models \text{dom}(a) \rightarrow v$. It follows that the same case is chosen in the definitions of $\mathbf{ta}_u^\diamond(\alpha a)$ and $\mathbf{ta}_u^\diamond(\beta a)$, and we obtain that $\mathbf{ta}_u^\diamond(\alpha a) = \mathbf{ta}_u^\diamond(\beta a)$.

It is immediate from the definition of \mathbf{ta}^\diamond that the relations $\{\sim_u^\diamond\}_{u \in D}$ satisfy DLR. By the previous paragraph, these relations also satisfy WSC. Since $\{\sim_u^{unw}\}_{u \in D}$ are the minimal relations satisfying WSC and DLR, we obtain that $\{\sim_u^{unw}\}_{u \in D} \leq \{\sim_u^\diamond\}_{u \in D}$. The converse also holds by Lemma 3. \square

From this, we obtain the claimed independence of locality on the choice of equivalence relation.

Lemma 5 *A given policy \rightarrow is local with respect to the associated relations $\{\sim_u^\diamond\}_{u \in D}$ iff it is local with respect to the relations $\{\sim_u^\square\}_{u \in D}$.*

Proof: First suppose that \succrightarrow is local with respect to $\{\sim_u^\diamond\}_{u \in D}$. Then by Lemma 4 and Theorem 3.1, we have $\{\sim_u^{unw}\}_{u \in D} = \{\sim_u^\diamond\}_{u \in D} = \{\sim_u^\square\}_{u \in D}$. Hence, \succrightarrow is local with respect to $\{\sim_u^\square\}_{u \in D}$.

Suppose now that \succrightarrow is local with respect to $\{\sim_u^\square\}_{u \in D}$. Let $u, v \in D$ and $\alpha, \beta \in A^*$ with $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$ and $\mathbf{ta}_v^\diamond(\alpha) = \mathbf{ta}_v^\diamond(\beta)$. By Lemma 3, we have $\alpha \sim_u^{unw} \beta$ and $\alpha \sim_v^{unw} \beta$. By Theorem 3.1, we have $\alpha \sim_u^\square \beta$ and $\alpha \sim_v^\square \beta$ and hence, $\alpha \models u \succrightarrow v$ iff $\beta \models u \succrightarrow v$. \square

Consequently, in the sequel we say simply that the policy is local, if it is local with respect to either of these relations.

The following result shows that for local policies, the conflict between a permissive and a prohibitive reading of policies does not arise.

Theorem 3.2 *If \succrightarrow is a local policy then \mathbf{ta}^\diamond -security and \mathbf{ta}^\square -security with respect to \succrightarrow are equivalent.*

Proof: This is immediate from the fact that if the policy is local, then for all traces α , we have $\alpha \models u \succrightarrow v$ iff $\alpha \models D_{\{u,v\}}(u \succrightarrow v)$. This claim follows straightforwardly from the properties of distributed knowledge. \square

Indeed, locality *exactly* captures the condition under which the permissive and the prohibitive interpretations of a policy are equivalent. We need one technical condition on this statement. Say that a domain u is *inactive* when $\text{dom}^{-1}(u) = \emptyset$, i.e., the domain has no actions that it can perform, so it plays only the role of an observer in the system. Note that policy edges of the form $\alpha \models u \succrightarrow v$ where u is inactive have no bearing on the definitions of \mathbf{ta}^\diamond -security or unwinding-security (and hence also not on the equivalent \mathbf{ta}^\square -security) since all references to the policy in these definitions occur only in the form $\alpha \models u \succrightarrow v$ with $u = \text{dom}(a)$ for some action a . Intuitively, the interference relation $u \succrightarrow v$ is concerned with the ability of domain u to perform actions that transfer information to domain v , and for a domain with no actions, such transfer of information is always impossible. We say that policy \succrightarrow *has no edges from inactive domains*, when for all domains u that are inactive, we have $\alpha \models u \not\succrightarrow v$ for all traces α and domains v . By the above observation, every policy is semantically equivalent to one that has no edges from inactive domains, so we may assume this condition without loss of generality. Subject to this assumption, we have the following:

Theorem 3.3 *Let $\mathcal{S} = \langle D, \text{dom}, A \rangle$ be a signature, and let \succrightarrow be a policy for this signature that has no edges from inactive domains. Suppose that for all systems M for the signature \mathcal{S} , M is \mathbf{ta}^\diamond -secure with respect to \succrightarrow iff M is \mathbf{ta}^\square -secure with respect to \succrightarrow . Then \succrightarrow is local.*

Proof: We prove the converse, that is, we show that if \succrightarrow is not local, then \mathbf{ta}^\diamond -security and \mathbf{ta}^\square -security with respect to \succrightarrow differ on some system M . Suppose that \succrightarrow is not local. By Lemma 5, we may use the relations $\{\sim_u^\diamond\}_{u \in D}$, so we have that there exist traces $\alpha, \beta \in A^*$ and domains u, v such that $\alpha \sim_u^\diamond \beta$ and $\alpha \sim_v^\diamond \beta$, and $\alpha \models u \succrightarrow v$ and $\beta \not\models u \succrightarrow v$.

By Lemma 3, we also have $\alpha \sim_u^{unw} \beta$ and $\alpha \sim_v^{unw} \beta$. Since \succrightarrow has no edges from inactive domains, and $\alpha \models u \succrightarrow v$, domain u has actions. Let a be any

action with $\text{dom}(a) = u$. Then, by condition WSC, it follows that $\alpha a \sim_v^{unw} \beta a$. Moreover, since $\alpha \sim_v^\diamond \beta$ and $\beta \sim_v^\diamond \beta a$ (because $\beta \models u \not\rightarrow v$), we have that $\alpha \sim_v^\diamond \beta a$. By $\alpha \models u \rightarrow v$ we have that $\alpha a \not\sim_v^\diamond \alpha$, so it follows that $\alpha a \not\sim_v^\diamond \beta a$.

Define M to be the system with states $S = A^*$, initial state $s_0 = \epsilon$, transitions defined by $\gamma \cdot b = \gamma b$ for all $\gamma \in A^*$ and $b \in A$, and observations defined by $\text{obs}_w(\gamma) = 0$ if $w \neq v$ or $\gamma \not\sim_v^\diamond \alpha a$, and $\text{obs}_w(\gamma) = 1$ otherwise, for all $\gamma \in A^*$. Obviously, $\text{obs}_v(\alpha a) = 1$. On the other hand, since $\alpha a \not\sim_v^\diamond \beta a$, we have that $\text{obs}_v(\beta a) = 0$.

Since $\alpha a \sim_v^{unw} \beta a$, it is immediate that M is not \mathbf{ta}^\square -secure with respect to \rightarrow . To complete the proof, we note that M is \mathbf{ta}^\diamond -secure with respect to \rightarrow , so that the two notions disagree on system M . To see this, note that any insecurity must involve domain v , since $\text{obs}_w(\gamma) = 0$ for all γ and $w \neq v$. However, on domain v , if $\gamma \sim_v^\diamond \gamma'$ then $\gamma \not\sim_v^\diamond \alpha a$ iff $\gamma' \not\sim_v^\diamond \alpha a$, so we have $\text{obs}_v(\gamma) = \text{obs}_v(\gamma')$, as required for \mathbf{ta}^\diamond -security. \square

Given this result, one reasonable approach to the possibility of conflicting interpretations of policies is to require that the policy be local, so that we are left with a single notion of security that supports both the permissive and prohibitive interpretation of policies. A potential disadvantage of this is that such a restriction results in a loss of expressiveness: certain policies can no longer be expressed. The following result shows that, in fact, provided that one is interested in \mathbf{ta}^\square -security, there is no loss of expressiveness: for every policy, there is a local policy that is equivalent with respect to \mathbf{ta}^\square -security.

Lemma 6 *For every policy \rightarrow over signature $\langle D, \text{dom}, A \rangle$, there is a local policy $\rightarrow' \leq \rightarrow$ such that for all systems M with signature $\langle D, \text{dom}, A \rangle$, system M is \mathbf{ta}^\square -secure with respect to \rightarrow iff M is \mathbf{ta}^\square -secure with respect to \rightarrow' .*

Proof: We write $[\alpha]_\sim$ for the equivalence class of α with respect to an equivalence relation \sim . In the following, the unwinding $\{\sim_u^{unw}\}_{u \in D}$ refers to the policy \rightarrow . Define the policy \rightarrow' by

$$\alpha \models u \rightarrow' v \text{ iff } \beta \models u \rightarrow v \text{ for all } \beta \in [\alpha]_{\sim_u^{unw}} \cap [\alpha]_{\sim_v^{unw}}.$$

It is clear that $\rightarrow' \leq \rightarrow$.

Let $\{\approx_u^{unw}\}_{u \in D}$ be the smallest unwinding satisfying WSC and DLR with respect to the policy \rightarrow' . We claim that $\sim_u^{unw} = \approx_u^{unw}$ for all $u \in D$. By Lemma 2, we have $\approx_u^{unw} \subseteq \sim_u^{unw}$ for all $u \in D$.

We will show that for every $u \in D$ and every $\alpha, \beta \in A^*$, we have $\alpha \approx_u^{unw} \beta$ implies $\alpha \sim_u^{unw} \beta$. The proof is by an induction on the length of the derivation of $\alpha \approx_u^{unw} \beta$. Consider a derivation of $\alpha \approx_u^{unw} \beta$ and suppose that the claim holds for all shorter derivations. We need to show that $\alpha \sim_u^{unw} \beta$. The cases where the final step of the derivation are an application of reflexivity, symmetry or transitivity are straightforward from the fact that \sim_u^{unw} is an equivalence relation. There are two remaining cases:

- *Case 1:* $\alpha \approx_u^{unw} \beta$ is derived using the condition DLR. Here we have that there exists $a \in A$ such that $\alpha = \beta a$ and $\beta \models \text{dom}(a) \not\rightarrow u$. Thus,

there exists a trace γ such that $\gamma \sim_u^{unw} \beta$ and $\gamma \sim_{\text{dom}(a)}^{unw} \beta$ such that $\beta \models \text{dom}(a) \not\rightarrow u$. By Proposition 1, we have that \sim_u^{unw} satisfies DLR'. It follows that $\alpha \sim_u^{unw} \beta$.

- *Case 2:* $\alpha \approx_u^{unw} \beta$ is derived using condition WSC. Hence we have $\alpha = \alpha'a$ and $\beta = \beta'a$ with $\alpha' \approx_u^{unw} \beta'$ and $\alpha' \approx_{\text{dom}(a)}^{unw} \beta'$. By applying the induction hypothesis, we have $\alpha' \sim_u^{unw} \beta'$ and $\alpha' \sim_{\text{dom}(a)}^{unw} \beta'$ and by the WSC condition applied to \sim_u^{unw} , we have $\alpha'a \sim_u^{unw} \beta'a$.

This shows that the system M is \mathbf{ta}^\square -secure with respect to \rightarrow if and only if it is \mathbf{ta}^\square -secure with respect to \rightarrow' .

It remains to show that the policy \rightarrow' is \mathbf{ta}^\diamond -local. First, we will show for every $u \in D$ and every $\alpha, \beta \in A^*$, that $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$ (with respect to \rightarrow') implies $\alpha \approx_u^{unw} \beta$. We prove this by an induction on the combined length of α and β . The base case of $\alpha = \beta = \epsilon$ is trivial by reflexivity of \approx_u^{unw} .

Assume the claim for all traces of combined length shorter than $|\alpha| + |\beta|$. Let $\alpha = \alpha'a$ and suppose that $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$. We need to show that $\alpha'a \approx_u^{unw} \beta$.

- *Case 1:* $\alpha' \models \text{dom}(a) \not\rightarrow' u$. In this case, we have $\mathbf{ta}_u^\diamond(\alpha'a) = \mathbf{ta}_u^\diamond(\alpha')$, so $\mathbf{ta}_u^\diamond(\alpha') = \mathbf{ta}_u^\diamond(\beta)$. By induction hypothesis, we have $\alpha' \approx_u^{unw} \beta$. By DLR, we have $\alpha'a \approx_u^{unw} \alpha'$, so by transitivity we obtain $\alpha'a \approx_u^{unw} \beta$.
- *Case 2:* $\alpha' \models \text{dom}(a) \rightarrow' u$. Then $\mathbf{ta}_u^\diamond(\alpha'a) = (\mathbf{ta}_u^\diamond(\alpha'), \mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha'), a)$, so $\beta \neq \epsilon$ and we have $\beta = \beta'b$ for some $b \in A$. We may assume that $\beta' \models \text{dom}(b) \rightarrow' u$, since otherwise we may argue as in Case 1 for $\beta'b$ and apply symmetry. It follows that $\mathbf{ta}_u^\diamond(\beta'b) = (\mathbf{ta}_u^\diamond(\beta'), \mathbf{ta}_{\text{dom}(b)}^\diamond(\beta'), b)$, and we conclude that $a = b$ and $\mathbf{ta}_u^\diamond(\alpha') = \mathbf{ta}_u^\diamond(\beta')$ and $\mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha') = \mathbf{ta}_{\text{dom}(a)}^\diamond(\beta')$. By induction hypothesis we have $\alpha' \approx_u^{unw} \beta'$ and $\alpha' \approx_{\text{dom}(a)}^{unw} \beta'$ and by WSC, we obtain that $\alpha'a \approx_u^{unw} \beta'a$.

The \mathbf{ta}^\diamond -locality follows: if $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$ and $\mathbf{ta}_v^\diamond(\alpha) = \mathbf{ta}_v^\diamond(\beta)$ (with respect to \rightarrow'), then we have $\alpha \in [\beta]_{\approx_u^{unw}} \cap [\beta]_{\approx_v^{unw}}$ and hence $\alpha \models u \rightarrow' v$ iff $\beta \models u \rightarrow' v$. \square

We say that two policies \rightarrow and \rightarrow' for a common signature $\langle D, \text{dom}, A \rangle$ are *identical up to inactive domains*, if for all domains $u, v \in D$ such that the set of actions $\text{dom}^{-1}(u)$ of domain u is nonempty, we have, for all $\alpha \in A^*$, that $\alpha \models u \rightarrow v$ iff $\alpha \models u \rightarrow' v$. Intuitively, this says that the two policies are identical, except that they may differ on edges $\alpha \models u \rightarrow' v$ from inactive domains u . Since such a domain u cannot take any action to cause information flow to v , this difference cannot, in practice, be detected. The following result makes this intuition precise. Say that two policies \rightarrow and \rightarrow' are two policies for a common signature are *equivalent* with respect to a notion of security if for all systems M for the signature, M is secure with respect to \rightarrow iff M is secure with respect to \rightarrow' .

Theorem 3.4 *Suppose that \rightarrow and \rightarrow' are two policies for a common signature $\langle D, \text{dom}, A \rangle$. Then \rightarrow and \rightarrow' are equivalent with respect to \mathbf{ta}^\diamond -security iff \rightarrow and \rightarrow' are identical up to inactive domains.*

Proof: In this proof, $\mathbf{ta}^{\diamond, \succ}$ and $\mathbf{ta}^{\diamond, \succ'}$ denote the \mathbf{ta}^{\diamond} function that refers to \succ and \succ' , respectively. Suppose first that \succ and \succ' are identical up to inactive domains. Then a straightforward induction on the length of $\alpha \in A^*$ shows that for all $u \in D$, we have $\mathbf{ta}_u^{\diamond, \succ}(\alpha) = \mathbf{ta}_u^{\diamond, \succ'}(\alpha)$. (Note that in the inductive case for $\mathbf{ta}_u^{\diamond}(\alpha a)$, we have that $a \in \text{dom}^{-1}(\text{dom}(u)) \neq \emptyset$, so $\text{dom}(a) \succ u$ iff $\text{dom}(a) \succ' u$.) It follows that \succ and \succ' are equivalent with respect to \mathbf{ta}^{\diamond} -security.

Conversely, suppose that \succ and \succ' are not identical up to inactive domains. We show that they are not equivalent with respect to \mathbf{ta}^{\diamond} -security. We have that there exists $\alpha \in A^*$ and domains u, v with $\text{dom}^{-1}(u)$ nonempty, such that (without loss of generality) $\alpha \models u \succ v$ but not $\alpha \models u \succ' v$. Let a be an action with $\text{dom}(a) = u$. Then $\mathbf{ta}_v^{\diamond, \succ}(\alpha a) \neq \mathbf{ta}_v^{\diamond, \succ}(\alpha)$ but $\mathbf{ta}_v^{\diamond, \succ'}(\alpha a) = \mathbf{ta}_v^{\diamond, \succ'}(\alpha)$. Define the system M with states $S = A^*$, transitions $\beta \cdot b = \beta b$ for all $\beta \in A^*$ and $b \in A$, and observations given by $\text{obs}_w(\beta) = 0$ if either $w \neq v$ or $\mathbf{ta}_w^{\diamond, \succ}(\beta) \neq \mathbf{ta}_w^{\diamond, \succ'}(\beta)$, otherwise $\text{obs}_w(\beta) = 1$. Then, by construction, M is \mathbf{ta}^{\diamond} -secure with respect to \succ . But we have $\text{obs}_v(\alpha a) = 0$ and $\text{obs}_v(\alpha) = 1$, so M is not \mathbf{ta}^{\diamond} -secure with respect to \succ' . Thus, these policies disagree on M . \square

In particular, it follows that if \succ' is a nonlocal policy that has no edges from inactive domains (so that the reason for nonlocality is not trivial, in the sense that it must involve edges from a domain with a nonempty set of actions), then there is no local policy that is equivalent with respect to \mathbf{ta}^{\diamond} -security. Thus, \mathbf{ta}^{\diamond} -security is inherently a more expressive notion than \mathbf{ta}^{\square} -security.

3.4 Special Cases

To judge the adequacy of a definition of security with respect to dynamic policies, it is useful to verify that the definition gives the expected answer in a number of simple scenarios where we have clear intuitions about the desired behaviour of the definition. The following provide some examples.

3.4.1 Static Policies:

We say that a dynamic policy $\succ \subseteq D \times A^* \times D$ is *static*, if for all $\alpha \in A^*$ and $u, v \in D$, we have $\alpha \models u \succ v$ iff $\epsilon \models u \succ v$. That is, the policy never changes from its initial state. A static policy \succ may equivalently be presented as the relation $\pi_{1,3}(\succ) \subseteq D \times D$, the projection of the dynamic policy to the first and third components, which is in precisely the format for classical intransitive noninterference policies [15].

Static policies are a special case of dynamic policies, that have been subject to a significant amount of study [17, 29, 30, 34], so it is reasonable to expect that any definition for dynamic policies should reduce to an accepted definition when applied to the static case. The following result shows that this is the case.

Theorem 3.5 *If \succ is a static policy, then the following are equivalent:*

- M is \mathbf{ta}^{\diamond} -secure with respect to \succ ,

- M is \mathbf{ta}^\square -secure with respect to \succrightarrow ,
- M is TA -secure with respect to $\pi_{1,3}(\succrightarrow)$.

Proof: This is immediate from the definitions, together with the fact that for static policies \succrightarrow , we have $\alpha \models u \succrightarrow v$ iff $\alpha \models D_{\{u,v\}}(u \succrightarrow v)$ iff $(u, v) \in \pi_{1,3}(\succrightarrow)$. \square

This result shows that \mathbf{ta}^\diamond -security and \mathbf{ta}^\square -security agree and collapse to the standard notion of \mathbf{ta} -security when restricted to static policies. This provides support for these definitions.

3.4.2 Globally known policies:

Another setting that can serve as a useful test case for the adequacy of definitions of security for dynamic policies is a group of agents subject to an information flow policy that is set by a policy authority P , and in which all agents are permitted to know the policy state at all times.

Formally, these are dynamic policies in which the set of domains contains P , and we have the following:

- (GK1) For all traces α and domains u , we have that $\alpha \models P \succrightarrow u$. That is, domain P , which controls the policy state, is always permitted to transmit information to u . This ensures that u is always able to know the policy state.
- (GK2) For all traces α, β and domains u, v , we have that $\alpha|P = \beta|P$ implies that $\alpha \models u \succrightarrow v$ iff $\beta \models u \succrightarrow v$. That is, the policy setting depends only on the past actions of domain P .

We say that a dynamic policy \succrightarrow satisfying these conditions is *globally known*.

Theorem 3.6 *If the policy \succrightarrow is globally known, then it is local.*

Proof: Suppose that the policy \succrightarrow is globally known. We need to show that for all $u, v \in D$ and $\alpha, \beta \in A^*$, we have $\alpha \sim_u^{unw} \beta$ and $\alpha \sim_v^{unw} \beta$ implies $\alpha \models u \succrightarrow v$ iff $\beta \models u \succrightarrow v$. We claim that, in fact, $\alpha \sim_u^{unw} \beta$ implies $\alpha|P = \beta|P$, from which the desired result follows using GK2.

The proof of the claim is by induction on the length of the derivation of $\alpha \sim_u^{unw} \beta$. The base case of a step in which $\alpha \sim_u^{unw} \beta$ is reflexivity by reflexivity is trivial, and the cases of steps using symmetry or transitivity are also straightforward.

Consider a step using DLR. Then $\alpha = \beta a$ and we have $\alpha \models \text{dom}(a) \succrightarrow u$. By GK1, we have $\text{dom}(a) \neq P$. Thus $\alpha|P = \beta a|P = \beta|P$, as required.

Finally, consider a step using WSC. Then $\alpha = \alpha' a$ and $\beta = \beta' a$ for some action a , and we have $\alpha' \sim_u^{unw} \beta'$ and $\alpha' \sim_{\text{dom}(a)}^{unw} \beta'$. From $\alpha' \sim_u^{unw} \beta'$, we obtain by induction that $\alpha|P = \beta|P$. It follows that $\alpha|P = \alpha' a|P = \beta' a|P = \beta|P$, as required. \square

Since by Theorem 3.2, for local policies the definitions of \mathbf{ta}^\diamond -security and \mathbf{ta}^\square -security are equivalent, these definitions are equivalent for all globally known policies, so these definitions lend mutual support to each other in this case.

3.4.3 Locally known policies:

A slightly more general case in which the notions of \mathbf{ta}^\diamond -security and \mathbf{ta}^\square -security coincide are policies in which agents are always aware of their incoming and/or outgoing policy edges. Formally, define say that the policy \rightarrow is *locally known to the sender* (with respect to equivalence relations $\{\sim_u^{unw}\}_{u \in D}$) if we have, for all $u, v \in D$, that $\alpha \sim_u^{unw} \beta$ implies $\alpha \models u \rightarrow v$ iff $\beta \models u \rightarrow v$. Similarly, say that the policy \rightarrow is *locally known to the receiver* if we have, for all $u, v \in D$, that $\alpha \sim_u^{unw} \beta$ implies $\alpha \models v \rightarrow u$ iff $\beta \models v \rightarrow u$. It is easily seen that if the policy is locally known to the sender and/or the receiver (with respect to either \sim^\square or \sim^\diamond), then it is local. Thus, this again is a case where the notions of \mathbf{ta}^\diamond -security and \mathbf{ta}^\square -security are equivalent.

3.5 Related Work using Automaton Semantics

To make the case for our definitions above, we now compare them to related prior work that is also stated in an automaton based semantics. (We defer discussion of related programming language-based work to Section 7.) To the best of our knowledge, the only other semantics for dynamic intransitive noninterference policies in automaton-based systems are those of Leslie [22] and Eggert et. al. [11]. Both works deal with policy enhanced systems. Recall that in this model, we write $s \models u \rightarrow v$ iff information flow from u to v is allowed in the state s .

Leslie's definition of intransitive noninterference is a direct adaption of Rushby's ipurge function [30] to the dynamic setting. Similar to Rushby's ipurge function, Leslie's purge function is based on a "source" function. Formally, the dynamic version of this function is defined for every $u \in D$ and $s \in S$ by $\mathbf{dsrc}(\epsilon, u, s) = \{u\}$ and for $\alpha \in A^*$ and $a \in A$ by $\mathbf{dsrc}(a\alpha, u, s) = \mathbf{dsrc}(\alpha, u, s \cdot a) \cup \{\mathbf{dom}(a)\}$ if there is some $v \in \mathbf{dsrc}(\alpha, u, s \cdot a)$ with $s \models \mathbf{dom}(a) \rightarrow v$, and else $\mathbf{dsrc}(a\alpha, u, s) = \mathbf{dsrc}(\alpha, u, s \cdot a)$. Intuitively, the set $\mathbf{dsrc}(a\alpha, u, s)$ is the set of all agents, to which the information that the action a has been performed, is transmitted by the sequence of actions α .

Leslie's purge function for intransitive noninterference is inductively defined for every $u \in D$, $s \in S$, $a \in A$, and $\alpha \in A^*$ by $\mathbf{Lpurge}(\epsilon, u, s) = \epsilon$ and

$$\mathbf{Lpurge}(a\alpha, u, s) = a \mathbf{Lpurge}(\alpha, u, s \cdot a)$$

if there is some $v \in \mathbf{dsrc}(a\alpha, u, s)$ with $s \models \mathbf{dom}(a) \rightarrow v$, and else $\mathbf{Lpurge}(a\alpha, u, s) = \mathbf{Lpurge}(\alpha, u, s \cdot a)$. This function removes exactly those actions from a trace that are from agents where interference to u is forbidden by the policy. The corresponding security definition is:




Figure 2: Non-monotony of **Lpurge**-security

Definition 5 A system M is **Lpurge**-secure with respect to a policy \succrightarrow if for all domains u and sequences of actions α holds: $\text{obs}_u(s_0 \cdot \text{Lpurge}(\alpha, u, s_0)) = \text{obs}_u(s_0 \cdot \alpha)$.

Our notion of \mathbf{ta}^\diamond -security is incomparable to **Lpurge**-security, in the sense that neither of them implies the other. On the one hand, on systems with a static policy, **Lpurge**-security is equivalent to Rushby's IP-security and hence, as shown in [34], strictly weaker than TA-security, which is the equivalent to \mathbf{ta}^\diamond -security on such systems. On the other hand, the following example presents a \mathbf{ta}^\diamond -secure, but not **Lpurge**-secure system.

Example 6: We consider again the system in Figure 1. As shown in Example 2, this system is \mathbf{ta}^\diamond -secure. However the system is not **Lpurge**-secure, since we have $\text{obs}_B(s_0 \cdot \text{Lpurge}(pa, B, s_0)) = \text{obs}_B(s_0 \cdot \text{Lpurge}(a, B, s_1)) = \text{obs}_B(s_0 \cdot a) = 0 \neq 1 = \text{obs}_B(s_0 \cdot pa)$. \square

Similar to \mathbf{ta}^\diamond -security, **Lpurge**-security is not monotonic with respect to the restrictiveness order on policies.

Example 7: The system in Figure 2 has two different configurations of the dynamic policy indicated by the dotted policy edge in state s_1 : in one configuration there is an edge from D to L and in the other, this edge is absent. First, consider the configuration where this edge exists. With respect to **Lpurge**-security, this system is insecure, since we have $\text{Lpurge}(hd, L, s_0) = \text{Lpurge}(d, L, s_1) = d$, but $\text{obs}_L(s_0 \cdot hd) = 0 \neq 1 = \text{obs}_L(s_0 \cdot d)$.

However, if we remove the edge from D to L in state s_1 , this system becomes **Lpurge**-secure: all traces starting with h will be purged to ϵ and all traces starting with d have a purge value different from ϵ . \square

Our notion of \mathbf{ta}^\square is strictly stronger than **Lpurge**-security. That these two are indeed different follows from the same argument as above for systems with static policies.

Lemma 7 Every \mathbf{ta}^\square -secure system is **Lpurge**-secure.

Proof: We show this lemma by contraposition. Suppose that the system is not

Lpurge-secure. Then there are $u \in D$ and $\alpha \in A^*$ with

$$\text{obs}_u(s_0 \cdot \alpha) \neq \text{obs}_u(s_0 \cdot \text{Lpurge}(\alpha, u, s_0)) .$$

We will show the following claim: For every $\beta, \beta' \in A^*$ with $\alpha = \beta\beta'$ and for every $\gamma \in A^*$ with $\text{Lpurge}(\alpha, u, s_0) = \gamma \text{Lpurge}(\beta', u, s_0 \cdot \beta)$ and for every $v \in \text{dsrc}(\beta', u, s_0 \cdot \beta)$, we have $\beta \sim_v^{unw} \gamma$.

We prove this claim by an induction on β . For the base case with $\beta = \epsilon$, we have that $\alpha = \beta'$ and hence $\gamma = \epsilon$ from which the claim is immediate. For the inductive step, let $\beta = \tilde{\beta}b$ for some action b . The induction hypothesis is: For some γ with $\text{Lpurge}(\alpha, u, s_0) = \gamma \text{Lpurge}(b\beta', u, s_0 \cdot \tilde{\beta})$ and every $v \in \text{dsrc}(b\beta', u, s_0 \cdot \tilde{\beta})$, we have $\tilde{\beta} \sim_v^{unw} \gamma$.

We consider the following two cases:

- *Case 1:* $\text{dom}(a) \notin \text{dsrc}(b\beta', u, s_0 \cdot \tilde{\beta})$. In this case, we have

$$\text{dsrc}(b\beta', u, s_0 \cdot \tilde{\beta}) = \text{dsrc}(\beta', u, s_0 \cdot \tilde{\beta}b)$$

and

$$\text{Lpurge}(b\beta', u, s_0 \cdot \tilde{\beta}) = \text{Lpurge}(\beta', u, s_0 \cdot \tilde{\beta}b) .$$

Hence, the new value of γ is the same as the previous value. And since $s_0 \cdot \tilde{\beta} \models \text{dom}(b) \mapsto v$, by DLR, for every $v \in \text{dsrc}(\beta', u, s_0 \cdot \tilde{\beta}b)$, we have $\tilde{\beta}b \sim_v^{unw} \tilde{\beta}$, and combined with $\tilde{\beta} \sim_v^{unw} \gamma$ the claim follows.

- *Case 2:* $\text{dom}(a) \in \text{dsrc}(b\beta', u, s_0 \cdot \tilde{\beta})$. In this case, we have $\text{dsrc}(b\beta', u, s_0 \cdot \tilde{\beta}) = \{\text{dom}(b)\} \cup \text{dsrc}(\beta', u, s_0 \cdot \tilde{\beta}b)$. Since

$$\begin{aligned} \text{Lpurge}(\alpha, u, s_0) &= \gamma \text{Lpurge}(b\beta', u, s_0 \cdot \tilde{\beta}) \\ &= \gamma b \text{Lpurge}(\beta', u, s_0 \cdot \tilde{\beta}b) , \end{aligned}$$

the new value of γ is $\gamma' = \gamma b$. For every $v \in \text{dsrc}(\beta', u, s_0 \cdot \tilde{\beta}b)$, we have $\tilde{\beta} \sim_v^{unw} \gamma$ by induction hypothesis and additionally $\tilde{\beta} \sim_{\text{dom}(b)}^{unw} \gamma$. From the condition WSC it follows for every such v : $\tilde{\beta}b \sim_v^{unw} \gamma b = \gamma'$. From this claim it follows $\alpha \sim_u^{unw} \text{Lpurge}(\alpha, u, s_0)$ with $\beta = \alpha$ and hence the system is *not* ta^\square -secure. \square

In [11], some problems with Lpurge-security are identified and a new purge-based security definition is proposed in response. This notion uses the same dynamic sources function as Leslie's does, but the purge function does not change in the state parameter if an action has been removed. More precisely, their dynamic intransitive purge function is defined for every $u \in D$ and $s \in S$ by $\text{ipurge}(\epsilon, u, s) = \epsilon$ and for every $a \in A$ and $\alpha \in A^*$ if $\text{dom}(a) \in \text{dsrc}(a\alpha, u, s)$, then $\text{ipurge}(a\alpha, u, s) = a \text{ipurge}(\alpha, u, s \cdot a)$ and else $\text{ipurge}(a\alpha, u, s) = \text{ipurge}(\alpha, u, s)$. The notion of *i-security* is defined as usual as the property that for all agent u , all states s , and all action sequences α and β with $\text{ipurge}(\alpha, u, s) = \text{ipurge}(\beta, u, s)$, we have $\text{obs}_u(s \cdot \alpha) = \text{obs}_u(s \cdot \beta)$.




Figure 3: A \mathbf{ta}^\square -secure, but *not* i-secure system




Figure 4: Relations between dynamic intransitive noninterference definitions

This notion of security is incomparable with our security definitions of \mathbf{ta}^\diamond -security and \mathbf{ta}^\square -security. On the one hand, we have, due to the static case, that i-security does not imply \mathbf{ta}^\diamond -security. The following example gives a system, with a local policy, that is \mathbf{ta}^\square -secure and \mathbf{ta}^\diamond -secure, but not i-secure.

Example 8: The system in Figure 3 is clearly *not* i-secure. We consider the traces hd and d since the \mathbf{ipurge} values of these two traces are the same for L if one starts purging in s_0 : $\mathbf{ipurge}(hd, L, s_0) = \mathbf{ipurge}(d, L, s_0) = d$. But the observations after these two traces are different for L .

However, for any trace that starts with an h action, the \mathbf{ta}^\diamond value is ϵ , and for any trace that starts with a d action, the \mathbf{ta}^\diamond value differs from ϵ . Hence, for any two traces that lead to different observations for L have different \mathbf{ta}^\diamond values. Note, that the dynamic policy of this system is local, and hence, this system is also \mathbf{ta}^\square -secure.

We would argue that this system is intuitively secure, so this gives grounds to prefer our two definitions over i-security. First, note that in every state D is permitted to learn about H actions, so, in particular, it should be permitted for D to know in state s_0 that H has not yet acted. In state s_2 , L can deduce that D acted before H did, but D was permitted to transfer its information to L when it acted, so this is not an insecurity. Finally, L cannot distinguish state s_1 from s_0 , so this state is secure on the grounds that the initial state of a system always is (noninterference definitions, intuitively, aim to prohibit agents from learning information about what has happened in a system). \square

Although i-security implies Lpurge-security, Example 8 shows that these security definitions still differ on systems with local policies. The implications of the security definitions analyzed in this work are depicted in Figure 4. All these implications are strict in general, and the definitions of \mathbf{ta}^\square -security and \mathbf{ta}^\diamond -security only collapse on systems with local policies.

We can summarize our reasons for preferring the definitions of the present paper to these prior works as follows. First, both **Lpurge**-security and **i**-security are equivalent to Rushby's IP-security in the case of static policies. This means that these definitions are subject to the problems identified by van der Meyden [34]. One of these problems is an example presented in [34], which shows that Rushby's IP-security fails to detect a covert channel based on ordering of actions. A second weakness is that IP-security has a weaker relationship to the unwinding proof method for intransitive non-interference proposed by Rushby [30] than does **ta**-security. As shown in [34], this proof method is sound but not complete for IP-security, but both sound and complete for TA-security. Since the proof method is itself based on well-founded intuitions, this lends further support to TA-security. (We establish a related completeness result for \mathbf{ta}^\square -security and \mathbf{ta}^\diamond -security with respect to local dynamic policies in the following section, that shows that a similar justification exists for these extensions of TA-security to the dynamic case.) Secondly, when we consider **Lpurge**-security and **i**-security with respect to dynamic policies, there are examples for each where they give intuitively incorrect results. We refer to [11] for examples showing **Lpurge**-security yields undesirable conclusions. Example 8 shows that **i**-security can also yield undesirable results.

3.6 Policy Refinement

Recall that $\succ \leq \succ'$ says that policy \succ is more restrictive than policy \succ' . We therefore expect that if a system is secure with respect to \succ , it is also secure with respect to \succ' . The following result shows that \mathbf{ta}^\square -security behaves as expected with respect to this order.

Proposition 2 *If $\succ \leq \succ'$ and system M is \mathbf{ta}^\square -secure with respect to \succ , then M is \mathbf{ta}^\square -secure with respect to \succ' .*

Proof: This follows from the equivalence of \mathbf{ta}^\square -security and unwinding-security. Let $\succ \leq \succ'$ and let $\{\sim_u^{unw}\}_{u \in D}$ and $\{\approx_u^{unw}\}_{u \in D}$ be the unwinding relations for the policies \succ and \succ' , respectively. By Lemma 2, we have $\approx_u^{unw} \subseteq \sim_u^{unw}$ for all $u \in D$. Suppose that M is unwinding secure with respect to \succ . Then $\alpha \approx_u^{unw} \beta$ implies $\alpha \sim_u^{unw} \beta$, which yields $\mathbf{obs}_u(s_0 \cdot \alpha) = \mathbf{obs}_u(s_0 \cdot \beta)$, as required for M to be secure with respect to \succ' . \square

The following example illustrates a further subtlety concerning \mathbf{ta}^\diamond -security: unlike \mathbf{ta}^\square -security, it is not monotonic with respect to the restrictiveness order on policies.

Example 9: Figure 5 shows a system and a policy \succ for two agents A, B . Only the observations of agent B are depicted; for agent A we assume that $\mathbf{obs}_A(s) = 0$ for all states s , so that there can be no insecurity with respect to agent A . For agent B , the only possible cases of \mathbf{ta}^\diamond -insecurity are when $\mathbf{obs}_B(s_0 \cdot \alpha) = 1$ and $\mathbf{obs}_B(s_0 \cdot \beta) = 0$. This implies that we have $\alpha = abaa'$ and that aba is not a prefix of β . However, in this case, $\mathbf{ta}_B^\diamond(\beta)$ cannot contain




Figure 5: The permissive interpretation is not monotonic: policy \rightsquigarrow

a subterm a , since there are no strict prefixes γ of β with $\gamma \models A \rightsquigarrow B$. On the other hand, $\mathbf{ta}_B^\diamond(\text{aba}\alpha')$ always contains a subterm a , since $\mathbf{ta}_B^\diamond(\text{aba}) = ((\epsilon, \epsilon, b), (\epsilon, \epsilon, a), a)$. Thus, we cannot have $\mathbf{ta}_B^\diamond(\text{aba}\alpha) = \mathbf{ta}_B^\diamond(\beta)$. It follows that the system is \mathbf{ta}^\diamond -secure with respect to \rightsquigarrow . We note that the policy \rightsquigarrow is not local, because $\mathbf{ta}_A^\diamond(ab) = \mathbf{ta}_A^\diamond(ba) = (\epsilon, \epsilon, a)$ and $\mathbf{ta}_B^\diamond(ab) = \mathbf{ta}_B^\diamond(ba) = (\epsilon, \epsilon, b)$, but $ab \models A \rightsquigarrow B$ while $ba \models A \not\rightsquigarrow B$.

On the other hand, consider the policy \rightsquigarrow' , which is defined identically to \rightsquigarrow , except that on state s_6 we take $s_6 \models A \rightsquigarrow' B$. Plainly, $\rightsquigarrow \leq \rightsquigarrow'$. However, the system is not \mathbf{ta}^\diamond -secure with respect to \rightsquigarrow' . Note that with respect to policy \rightsquigarrow' , we have $\mathbf{ta}_B^\diamond(\text{aba}) = ((\epsilon, \epsilon, b), (\epsilon, \epsilon, a), a) = \mathbf{ta}_B^\diamond(\text{baa})$, whereas $\text{obs}_B(s_0 \cdot \text{aba}) = 1$ and $\text{obs}_B(s_0 \cdot \text{baa}) = 0$. \square

However, for local policies, \mathbf{ta}^\diamond -security behaves as expected. This follows from Proposition 2, the equivalence of \mathbf{ta}^\diamond -security and \mathbf{ta}^\square -security for local policies. The following result shows that, in fact it holds slightly more generally, and we can also use monotonicity to derive \mathbf{ta}^\diamond -security for non-local policies.

Proposition 3 *Suppose that $\rightsquigarrow \leq \rightsquigarrow'$ and that \rightsquigarrow is local. If a system M is \mathbf{ta}^\diamond -secure with respect to \rightsquigarrow , then M is \mathbf{ta}^\diamond -secure with respect to \rightsquigarrow' .*

Proof: Let M be a \mathbf{ta}^\diamond -secure system with respect to \rightsquigarrow . Since \rightsquigarrow is local, we obtain from Theorem 3.2 that M is \mathbf{ta}^\square -secure with respect to \rightsquigarrow . By Proposition 2, the system is \mathbf{ta}^\square -secure with respect to \rightsquigarrow' . Since \mathbf{ta}^\square -security implies \mathbf{ta}^\diamond -security (by Lemma 3 and Theorem 3.1), we have that M is \mathbf{ta}^\diamond -secure with respect to \rightsquigarrow' . \square

4 Proof Techniques

In the theory of static (intransitive) noninterference policies, *unwinding relations* provide a proof technique that can be used to show that a system is secure. An unwinding relation is a relation on the states of the system that satisfies certain conditions. We show in this section that a similar technique applies to dynamic policies. In fact, the conditions that we need are very similar to those defined on the policy above, with the main difference being that they are stated with respect to states of the system rather than traces.

4.1 Proof technique for \mathbf{ta}^\square -security

Let M be a system and let \approx_u be an equivalence relation on the states of M for each $u \in D$.

(DLR $_M$) If $\alpha \models \text{dom}(a) \not\rightarrow u$ then $s_0 \cdot \alpha a \approx_u s_0 \cdot \alpha$.

(WSC $_M$) If $s \approx_u t$ and $s \approx_{\text{dom}(a)} t$ then $s \cdot a \approx_u t \cdot a$.

(OC $_M$) If $s \approx_u t$ then $\text{obs}_u(s) = \text{obs}_u(t)$.

We call a collection $\{\approx_u\}_{u \in D}$ of equivalence relations satisfying DLR $_M$, WSC $_M$ and OC $_M$ a \square -unwinding on M with respect to \rightarrow .

We remark that for policy-enhanced systems, we have the following version of DLR $_M$:

(DLR-PE $_M$) If $s \models \text{dom}(a) \not\rightarrow u$ then $s \cdot a \approx_u s$.

which formulates the conditions to uniformly quantify over states of the system rather than both over states and traces. It is straightforward to show that if a policy-enhanced system M satisfies DLR-PE $_M$ then it satisfies DLR $_M$.

The following result states that to prove \mathbf{ta}^\square -security, it suffices to show that there exists a \square -unwinding on the system.

Theorem 4.1 *Suppose that there exist equivalence relations $\{\approx_u\}_{u \in D}$ on the states of a system M satisfying DLR $_M$, WSC $_M$ and OC $_M$ with respect to a policy \rightarrow . Then M is \mathbf{ta}^\square -secure with respect to \rightarrow .*

Proof: To show that M is \mathbf{ta}^\square -secure, we argue that the relations $\{\sim_u^{unw}\}_{u \in D}$, which by definition satisfy WSC and DLR, also satisfy OC. For this, we claim that for all $\alpha, \beta \in A^*$, if $\alpha \sim_u^{unw} \beta$ then $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$. It is then immediate from OC $_M$ that OC holds for \sim_u^{unw} .

The proof of the claim is by induction on the derivation of $\alpha \sim_u^{unw} \beta$. The cases of $\alpha \sim_u^{unw} \beta$ obtained by reflexivity, symmetry or transitivity are straightforward, using the fact that \approx_u has these properties. For DLR, if $\alpha \sim_u^{unw} \beta$ is obtained by an application of DLR then $\alpha = \beta a$ for some $a \in A$ with $\alpha \models \text{dom}(a) \not\rightarrow u$. By DLR $_M$, we have $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$, as required. For WSC, suppose that $\alpha \sim_u^{unw} \beta$ is obtained by an application of WSC, so that we have $\alpha = \alpha' a$ and $\beta = \beta' a$, where $\alpha \sim_u^{unw} \beta$ and $\alpha \sim_{\text{dom}(a)}^{unw} \beta$. By induction, we have $s_0 \cdot \alpha' \approx_u s_0 \cdot \beta'$ and $s_0 \cdot \alpha' \approx_{\text{dom}(a)} s_0 \cdot \beta'$. We now obtain $s_0 \cdot \alpha' a \approx_u s_0 \cdot \beta' a$ by WSC $_M$, i.e., $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$, as required. \square

An unwinding on the system can often be naturally defined. We give an example below in the case of access control systems.

We may also show that unwindings provide a complete proof method, modulo the fact that we may need to work in a bisimilar system, rather than the system as presented. (See [34] for an example that demonstrates that, already for static policies, we do not obtain completeness if we require that the unwinding be defined over the states of the system M itself. A more complex type of

unwinding on the system M itself is shown to be complete for TA-security of static policies in [12]. It would be interesting to develop a generalization of this to the dynamic case, but we leave this for future work.)

Theorem 4.2 *Suppose that M is \mathbf{ta}^\square -secure with respect to \mapsto . Then there exists a system M' that is bisimilar to M (specifically, we may take $M' = \text{unfold}(M)$), such that there exist equivalence relations $\{\approx'_u\}_{u \in D}$ on the states of M' satisfying $\text{DLR}_{M'}$, $\text{WSC}_{M'}$ and $\text{OC}_{M'}$ with respect to policy \mapsto .*

Proof: Suppose that M is \mathbf{ta}^\square -secure with respect to \mapsto . Then the relations $\{\sim_u^{unw}\}_{u \in D}$ are defined on the states of $M' = \text{unfold}(M)$, and by definition, satisfy WSC and DLR . These are equivalent to $\text{WSC}_{M'}$ and $\text{DLR}_{M'}$. Moreover, by Theorem 3.1, we have that $\sim_u^{unw} = \sim_u^\square$ for all $u \in D$, so $\text{OC}_{M'}$ follows from the fact that M is \mathbf{ta}^\square -secure. \square

4.2 Proof technique for \mathbf{ta}^\diamond -security

A similar proof theory can be developed for the notion \mathbf{ta}^\diamond -security. We replace the condition WSC_M by the following:

(WSC_M^\diamond) If $\alpha \models \text{dom}(a) \mapsto u$ and $\beta \models \text{dom}(a) \mapsto u$ and $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$ and $s_0 \cdot \alpha \approx_{\text{dom}(a)} s_0 \cdot \beta$ then $s_0 \cdot \alpha a \approx_u s_0 \cdot \beta a$.

On policy enhanced systems M , it is easily seen that the following variant is sufficient for WSC_M^\diamond to hold:

(WSC-PE_M^\diamond) If $s \models \text{dom}(a) \mapsto u$ and $t \models \text{dom}(a) \mapsto u$ and $s \approx_u t$ and $s \approx_{\text{dom}(a)} t$ then $s \cdot a \approx_u t \cdot a$.

Theorem 4.3 *Suppose that there exist equivalence relations $\{\approx_u\}_{u \in D}$ on the states of a system M satisfying DLR_M , WSC_M^\diamond and OC_M with respect to a policy \mapsto . Then M is \mathbf{ta}^\diamond -secure with respect to \mapsto .*

Proof: We claim that for all $\alpha, \beta \in A^*$ we have $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$ implies that $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$. It follows from this using OC_M that M is \mathbf{ta}^\diamond -secure.

To prove the claim we proceed by induction on $|\alpha| + |\beta|$. The base case of $\alpha = \beta = \epsilon$ is trivial. For the inductive case, suppose (without loss of generality) that $\alpha = \alpha'a$, and we have $\mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$. There are two possibilities:

- $\alpha \models \text{dom}(a) \not\mapsto u$. In this case we have $\mathbf{ta}_u^\diamond(\alpha') = \mathbf{ta}_u^\diamond(\alpha) = \mathbf{ta}_u^\diamond(\beta)$, so by the inductive hypothesis, we have $s_0 \cdot \alpha' \approx_u s_0 \cdot \beta$. By DLR_M , we also have that $s_0 \cdot \alpha'a \approx_u s_0 \cdot \alpha'$. Since \approx_u is an equivalence relation, we have $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$.
- $\alpha \models \text{dom}(a) \mapsto u$. In this case, $\mathbf{ta}_u^\diamond(\alpha) = (\mathbf{ta}_u^\diamond(\alpha), \mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha), a)$. It follows that $\beta \neq \epsilon$, so we may write $\beta = \beta'b$ for some $\beta' \in A^*$ and $b \in A$. If $\beta' \models \text{dom}(b) \not\mapsto u$, we may reverse the roles of α and β and obtain that $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$ using the argument of the previous case. If $\beta' \models$

$\text{dom}(b) \mapsto u$, then $\text{ta}_u^\diamond(\beta) = (\text{ta}_u^\diamond(\beta'), \text{ta}_{\text{dom}(b)}^\diamond(\beta'), b)$, and we conclude that $a = b$ and $\text{ta}_u^\diamond(\alpha') = \text{ta}_u^\diamond(\beta')$ and $\text{ta}_{\text{dom}(a)}^\diamond(\alpha') = \text{ta}_{\text{dom}(a)}^\diamond(\beta')$. By the inductive hypothesis, we have $s_0 \cdot \alpha' \approx_u s_0 \cdot \beta'$ and $s_0 \cdot \alpha' \approx_{\text{dom}(a)} s_0 \cdot \beta'$. Using WSC_M^\diamond it follows that $s_0 \cdot \alpha' a \approx_u s_0 \cdot \beta' a$, i.e., $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$. \square

We also have completeness of the technique, provided we allow its application on a bisimilar system:

Theorem 4.4 *Suppose that M is ta^\square -secure with respect to \mapsto . Then there exists a system M' that is bisimilar to M and there exist equivalence relations $\{\approx_u\}_{u \in D}$ on the states of a system M' satisfying $\text{DLR}_{M'}$, $\text{WSC}_{M'}^\diamond$ and $\text{OC}_{M'}$ with respect to policy \mapsto .*

Proof: Let $M = \langle S, s_0, \rightarrow, D, \text{dom}, A, \text{obs} \rangle$. We take $M' = \langle S', s'_0, \rightarrow', D, \text{dom}, A, \text{obs}' \rangle$ to be the system over the same signature (D, dom, A) with states $S' = A^*$, initial state $s'_0 = \epsilon$, transitions \rightarrow' defined by $\alpha \cdot a = \alpha a$ for all $\alpha \in A^*$, and observations defined $\text{obs}'_u(\alpha) = \text{obs}_u(s_0 \cdot \alpha)$ for all $\alpha \in S'$. It is easily checked that this system is bisimilar to M .

We show that the relations $\{\sim_u^\diamond\}_{u \in D}$ satisfy $\text{DLR}_{M'}$, $\text{WSC}_{M'}^\diamond$ and $\text{OC}_{M'}$ with respect to policy \mapsto . Note that for all $\alpha \in A^*$, we have $s'_0 \cdot \alpha = \alpha$. For condition $\text{DLR}_{M'}$, note that if $\alpha \models \text{dom}(a) \not\mapsto u$ then $\text{ta}_u^\diamond(\alpha a) = \text{ta}_u^\diamond(\alpha)$, so $\alpha a \sim_u^\diamond \alpha$. For condition $\text{WSC}_{M'}^\diamond$, note that if $\alpha \models \text{dom}(a) \not\mapsto u$ and $\beta \models \text{dom}(a) \not\mapsto u$ and $\alpha \sim_u^\diamond \beta$ and $\alpha \sim_{\text{dom}(a)}^\diamond \beta$ then

$$\begin{aligned} \text{ta}_u^\diamond(\alpha a) &= (\text{ta}_u^\diamond(\alpha), \text{ta}_{\text{dom}(a)}^\diamond(\alpha), a) \\ &= (\text{ta}_u^\diamond(\beta), \text{ta}_{\text{dom}(a)}^\diamond(\beta), a) \\ &= \text{ta}_u^\diamond(\beta a) \end{aligned}$$

hence $\alpha a \sim_u^\diamond \beta a$. Condition $\text{OC}_{M'}$ is immediate from the fact that M (and hence M') is ta^\diamond -secure. \square

5 Access Control in Dynamic Systems

Access control monitors are a fundamental security mechanism that can be used to enforce a variety of policies. Rushby [30] showed that access control can provide a sound enforcement mechanism for static intransitive noninterference policies, and that a simple static check of the access control table suffices to verify the policy will be enforced. It was later shown by van der Meyden [34] that, in fact, access control is also a *complete* enforcement mechanism for such policies, in the sense that every system that satisfies the policy is bisimilar to one using access control, with a setting that satisfies the static check. In this section, we develop a generalization of these results to dynamic noninterference policies.

A system with structured state, is a system equipped with

- a set of all possible objects O ,

- a set of all possible values V ,
- a function **contents**: $O \times S \rightarrow V$,
- a function **observe**: $D \times S \rightarrow \mathcal{P}(O)$,
- a function **alter**: $D \times S \rightarrow \mathcal{P}(O)$.

Intuitively, **contents**(o, s) is the value of object o in state s , **observe**(u, s) is the set of objects that domain u is permitted to observe (or read) in state s , and **alter**(u, s) is the set of objects that domain u is permitted to alter (or write) in state s . The pair of functions **observe**, **alter** can be thought of as an encoding of an access control table that varies with the state of the system.

We make an assumption on the set O of objects, namely, that this contains a special object **oset**(u) for each domain $u \in D$. We write the set of all these special objects as $\mathbf{Osets}(D) = \{\mathbf{oset}(u) \mid u \in D\}$ and assume $\mathbf{Osets}(D) \subseteq O$. We assume that **contents**(**oset**(u), s) = **observe**(u, s), so that the content of **oset**(u) is the set of all objects that are observable to u . Furthermore, we assume that **oset**(u) \in **observe**(u, s) for all domains $u \in D$ and states $s \in S$, i.e., the object **oset**(u) is always observable by u . This is intuitively reasonable, given the meaning of this object, since domain u can always determine the value of **oset**(u) by checking directly which objects it is able to observe.

The motivation for these special objects **oset**(u) is that they are a kind of “meta-object” that can be used to specify if an agent is allowed to change the observable objects of another agent. More precisely, we may let **oset**(u) \in **alter**(v, s) if and only if in state s , domain v is allowed to influence which objects are observable by u .

We now state a number of conditions (resembling Rushby’s “reference monitor conditions” [30]) that are intended to capture the intuitive semantics of the access control functions **observe**, **alter**, as they might be enforced by a reference monitor that restricts agents’ ability to read and write objects as these access control settings change. The conditions are cast in terms of the relations \approx_u , for $u \in D$, defined by $s \approx_u t$ iff for all $o \in \mathbf{observe}(u, s)$, we have **contents**(o, s) = **contents**(o, t). Intuitively, $s \approx_u t$ when the objects observable by domain u have the same values in the states s and t . Note that because **oset**(u) \in **observe**(u, s), it is always the case that **observe**(u, s) = **observe**(u, t) when $s \approx_u t$. It follows that the relations \approx_u are in fact equivalence relations. Our new “dynamic reference monitor” conditions are the following:

DRM-1 If $s \approx_u t$, then **obs** _{u} (s) = **obs** _{u} (t).

DRM-2 If $o \in \mathbf{alter}(\mathbf{dom}(a), s) \cap \mathbf{alter}(\mathbf{dom}(a), t)$ and $s \approx_{\mathbf{dom}(a)} t$ and **contents**(o, s) = **contents**(o, t) then **contents**($o, s \cdot a$) = **contents**($o, t \cdot a$).

DRM-3 If **contents**($o, s \cdot a$) \neq **contents**(o, s), then $o \in \mathbf{alter}(\mathbf{dom}(a), s)$.

DRM-4 $\mathbf{observe}(u, s \cdot a) \setminus \mathbf{observe}(u, s) \subseteq \mathbf{observe}(\mathbf{dom}(a), s)$.

DRM-5 If $\alpha \models v \rightarrow u$ and $\beta \models v \rightarrow u$ and $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$ and $s_0 \cdot \alpha \approx_v s_0 \cdot \beta$, then $\mathbf{observe}(u, s_0 \cdot \alpha) \cap \mathbf{alter}(v, s_0 \cdot \alpha) = \mathbf{observe}(u, s_0 \cdot \beta) \cap \mathbf{alter}(v, s_0 \cdot \beta)$.

DRM-6 If $\mathbf{alter}(u, s_0 \cdot \alpha) \cap \mathbf{observe}(v, s_0 \cdot \alpha) \neq \emptyset$, then $\alpha \models u \rightarrow v$.

Intuitively, condition DRM-1 says that a domain's observation depends only on the contents of the objects observable to the domain. Condition DRM-2 says that when a domain is permitted to alter an object, the new value of the object depends only on its previous value and the values of the objects observable by the acting domain. (The version of this constraint given here follows [34] rather than [30].) Condition DRM-3 can be understood as stating that an action can change the value of an object only when the domain of the action is permitted by the access control setting to alter the object. Condition DRM-4 states that any object that is made newly observable to a domain u by an action a must be observable to the domain $\mathbf{dom}(a)$.

Condition DRM-5 is a kind of locality constraint, which states that when v is permitted to interfere with u , in situation s , the set $\mathbf{observe}(u, s) \cap \mathbf{alter}(v, s)$ of objects that u can observe and v can alter depends only on information local to domains u and v , as captured by the relations \approx_u and \approx_v . We note that this constraint actually only states a weak property of $\mathbf{alter}(v, s)$, as we already have that $\mathbf{observe}(u, s)$ depends only on the local state of u as captured by \approx_u . The following stronger version of this condition is also of interest:

DRM-5' If $s \approx_u t$ and $s \approx_v t$, then $\mathbf{observe}(u, s) \cap \mathbf{alter}(v, s) = \mathbf{observe}(u, t) \cap \mathbf{alter}(v, t)$.

Here we have dropped the requirement that v be permitted to interfere with u in order for the dependency to hold. This stronger condition more closely resembles locality.

We remark that have not introduced special objects $\mathbf{aset}(u)$ for each domain u , with the property that $\mathbf{contents}(\mathbf{aset}(u), s) = \mathbf{alter}(u, s)$, as might be suggested by analogy with the objects $\mathbf{oset}(u)$. The results to follow do not require the introduction of these objects. However, in systems with such objects, and also satisfying the property $\mathbf{aset}(u) \in \mathbf{observe}(u, s)$ condition DRM-5' would be necessarily satisfied.

Finally, condition DRM-6 relates the access control setting to the policy. Note that if there exists an object that domain u is able to write and domain v is able to read, then this object provides an obvious channel for information to flow from u to v . The condition states that in this situation, the resulting channel for information flow is permitted by the information flow policy \rightarrow .

Theorem 5.1 *Suppose M is a system with structured state that satisfies properties DRM-1 to DRM-6 with respect to policy \rightarrow . Then M is \mathbf{ta}^\diamond -secure with respect to \rightarrow . If M also satisfies the stronger condition DRM-5', then M is \mathbf{ta}^\square -secure with respect to \rightarrow .*

Proof: Assume that M is a system with structured state that satisfies properties DRM-1 to DRM-6 with respect to policy \mapsto . We first show that M satisfies DLR_M and WSC_M^\diamond and OC_M with respect to the relations $\{\approx_u\}_{u \in D}$.

For property DLR_M , suppose that $\alpha \models \text{dom}(a) \not\mapsto u$. Then by DRM-6, we have $\text{alter}(\text{dom}(a), s_0 \cdot \alpha) \cap \text{observe}(u, s_0 \cdot \alpha) = \emptyset$. Thus, for all $o \in \text{observe}(u, s_0 \cdot \alpha)$, we have $o \notin \text{alter}(\text{dom}(a), s_0 \cdot \alpha)$, and by DRM-3 we obtain that $\text{contents}(o, s_0 \cdot \alpha) = \text{contents}(o, s_0 \cdot \alpha a)$. This shows that $s_0 \cdot \alpha a \approx_u s_0 \cdot \alpha$, as required.

For WSC_M^\diamond , suppose that $\alpha \models \text{dom}(a) \mapsto u$ and $\beta \models \text{dom}(a) \mapsto u$ and $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$ and $s_0 \cdot \alpha \approx_{\text{dom}(a)} s_0 \cdot \beta$. We need to show that $s_0 \cdot \alpha a \approx_u s_0 \cdot \beta a$, i.e., that for all $o \in \text{observe}(u, s_0 \cdot \alpha a)$, we have $\text{contents}(o, s_0 \cdot \alpha a) = \text{contents}(o, s_0 \cdot \beta a)$. Let $o \in \text{observe}(u, s_0 \cdot \alpha a)$. By DRM-4 we have that $o \in \text{observe}(u, s_0 \cdot \alpha) \cup \text{observe}(\text{dom}(a), s_0 \cdot \alpha)$. From $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$ and $s_0 \cdot \alpha \approx_{\text{dom}(a)} s_0 \cdot \beta$ we obtain that $\text{observe}(u, s_0 \cdot \alpha) = \text{observe}(u, s_0 \cdot \beta)$ and $\text{observe}(\text{dom}(a), s_0 \cdot \alpha) = \text{observe}(\text{dom}(a), s_0 \cdot \beta)$, and for all $o \in \text{observe}(u, s_0 \cdot \alpha) \cup \text{observe}(\text{dom}(a), s_0 \cdot \alpha)$ we have $\text{contents}(o, s_0 \cdot \alpha) = \text{contents}(o, s_0 \cdot \beta)$. By two applications of DRM-5 (one for the pair of domains $u, \text{dom}(a)$, and the other for the pair $\text{dom}(a), \text{dom}(a)$), we obtain that $o \in \text{alter}(\text{dom}(a), s_0 \cdot \alpha)$ iff $o \in \text{alter}(\text{dom}(a), s_0 \cdot \beta)$. It therefore suffices to consider two cases:

- Suppose $o \in \text{alter}(\text{dom}(a), s_0 \cdot \alpha) \cap \text{alter}(\text{dom}(a), s_0 \cdot \beta)$. Then by DRM-2 we get $\text{contents}(o, s_0 \cdot \alpha a) = \text{contents}(o, s_0 \cdot \beta a)$.
- Suppose $o \notin \text{alter}(\text{dom}(a), s_0 \cdot \alpha)$ and $o \notin \text{alter}(\text{dom}(a), s_0 \cdot \beta)$. Then by DRM-3 we have $\text{contents}(o, s_0 \cdot \alpha a) = \text{contents}(o, s_0 \cdot \alpha) = \text{contents}(o, s_0 \cdot \beta) = \text{contents}(o, s_0 \cdot \beta a)$.

The property OC_M is trivial, since this is just DRM-1. We now have that M satisfies DLR_M and WSC_M^\diamond and OC_M . By Theorem 4.3, it follows that M is ta^\diamond -secure with respect to \mapsto .

Suppose that M also satisfies the stronger condition DRM-5'. We show that M is ta^\square -secure with respect to \mapsto . Consider the policy \mapsto^M defined by $\alpha \models u \mapsto^M v$ if $\text{alter}(u, s_0 \cdot \alpha) \cap \text{observe}(v, s_0 \cdot \alpha) \neq \emptyset$. It is easily checked that M satisfies DRM-1 to DRM-6 with respect to policy \mapsto^M , so by the above, we conclude that M is ta^\diamond -secure with respect to \mapsto^M . We claim that the policy \mapsto^M is local. This being the case, we obtain using Theorem 3.2 that M is ta^\square -secure with respect to \mapsto^M . By DRM-6, we have that $\mapsto^M \leq \mapsto$, so by Proposition 2 we get that M is ta^\square -secure with respect to \mapsto .

To show that \mapsto^M is local, define the equivalence relations $\{\sim_u\}_{u \in D}$ on A^* by $\alpha \sim_u \beta$ iff $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$. It follows from the fact that M satisfies DLR_M with respect to $\{\approx_u\}_{u \in D}$ that M satisfies DLR with respect to $\{\sim_u\}_{u \in D}$. An argument similar to that above for WSC_M^\diamond , but using DRM-5' in place of DRM-5, shows that M also satisfies WSC with respect to $\{\sim_u\}_{u \in D}$. Thus, for the smallest equivalence relations $\{\sim_u^{unw}\}_{u \in D}$ satisfying WSC and DLR with respect to \mapsto^M , we have $\sim_u^{unw} \subseteq \sim_u$ for all $u \in D$. Suppose that $\text{ta}_u^\diamond(\alpha) = \text{ta}_u^\diamond(\beta)$ and $\text{ta}_v^\diamond(\alpha) = \text{ta}_v^\diamond(\beta)$ (with respect to policy \mapsto^M). We need to show that $\alpha \models u \mapsto^M v$ iff $\beta \models u \mapsto^M v$. By Lemma 3, we have that $\alpha \sim_u^{unw} \beta$ and $\alpha \sim_v^{unw} \beta$. By what was shown above, we obtain that $\alpha \sim_u \beta$ and $\alpha \sim_v \beta$.

β , so $s_0 \cdot \alpha \approx_u s_0 \cdot \beta$ and $s_0 \cdot \alpha \approx_v s_0 \cdot \beta$. By DRM-5' we now obtain that $\text{observe}(u, s_0 \cdot \alpha) \cap \text{alter}(v, s_0 \cdot \alpha) = \text{observe}(u, s_0 \cdot \beta) \cap \text{alter}(v, s_0 \cdot \beta)$. It follows that $\alpha \models u \rightsquigarrow^M v$ iff $\beta \models u \rightsquigarrow^M v$, as required. \square

Theorem 5.1 shows that access control mechanisms provide an implementation method that guarantees that a system constructed using this method is secure. It states a soundness result: any system that satisfies the reference monitor conditions with respect to a policy is secure for that policy.

It is also possible to prove a converse to this result, stating that any system that is secure with respect to a policy could have been constructed using access control mechanisms so as to satisfy the policy. To obtain such a result, we need to allow consideration of a bisimilar system. In particular, we focus on the *unfolding* of a system M .

Say that a system *admits an access control interpretation consistent with a policy* \rightsquigarrow , if it can be extended into a system with structured state by adding definitions of $O, V, \text{contents}, \text{observe}$ and alter in such a way as to satisfy conditions DRM-1 to DRM-6.

Theorem 5.2 *Let system M be \mathbf{ta}^\diamond -secure with respect to \rightsquigarrow . Then $\text{unfold}(M)$ admits an access control interpretation consistent with \rightsquigarrow .*

If M is \mathbf{ta}^\square -secure with respect to \rightsquigarrow , then there exists a local policy $\rightsquigarrow' \leq \rightsquigarrow$ such that $\text{unfold}(M)$ admits an access control interpretation consistent with \rightsquigarrow' , that additionally satisfies DRM-5'.

Proof: Suppose first that M is \mathbf{ta}^\diamond -secure with respect to \rightsquigarrow . Extend $\text{unfold}(M)$ to a system with structured state by the following definitions:

$$\begin{aligned} O &= D \cup \mathbf{Osets}(D) \\ \text{observe}(u, \alpha) &= \{u, \mathbf{oset}(u)\} \\ \text{alter}(u, \alpha) &= \{v \in D \mid \alpha \models u \rightsquigarrow v\} \\ \text{contents}(x, \alpha) &= \begin{cases} \mathbf{ta}_x^\diamond(\alpha) & \text{if } x \in D \\ \{u, \mathbf{oset}(u)\} & \text{if } x = \mathbf{oset}(u) \end{cases} \end{aligned}$$

We will show that the dynamic reference monitor conditions DRM-1 to DRM-6 hold.

DRM-1: Let $\alpha \approx_u \beta$. Then we have $u \in \text{observe}(u, \alpha) = \text{observe}(u, \beta)$ and $\mathbf{ta}_u^\diamond(\alpha) = \text{contents}(u, \alpha) = \text{contents}(u, \beta) = \mathbf{ta}_u^\diamond(\beta)$. It now follows by \mathbf{ta}^\diamond -security of M that $\text{obs}_u(\alpha) = \text{obs}_u^M(s_0 \cdot \alpha) = \text{obs}_u^M(s_0 \cdot \beta) = \text{obs}_u(\beta)$.

DRM-2: Let $o \in \text{alter}(\text{dom}(a), \alpha) \cap \text{alter}(\text{dom}(a), \beta)$ and $\alpha \approx_{\text{dom}(a)} \beta$ and $\text{contents}(o, \alpha) = \text{contents}(o, \beta)$. From $\alpha \approx_{\text{dom}(a)} \beta$ and the fact that, by definition, $\text{dom}(a) \in \text{observe}(\text{dom}(a), \alpha)$ we obtain $\mathbf{ta}_{\text{dom}(a)}^\diamond(\alpha) = \mathbf{ta}_{\text{dom}(a)}^\diamond(\beta)$. From $o \in \text{alter}(\text{dom}(a), \alpha) \cap \text{alter}(\text{dom}(a), \beta)$ it follows that $o \in D$ and $\alpha \models \text{dom}(a) \rightsquigarrow o$ and $\beta \models \text{dom}(a) \rightsquigarrow o$. From $\text{contents}(o, \alpha) = \text{contents}(o, \beta)$ it follows that $\mathbf{ta}_o^\diamond(\alpha) = \mathbf{ta}_o^\diamond(\beta)$. By the definition of the \mathbf{ta}^\diamond -operator, we obtain $\mathbf{ta}_o^\diamond(\alpha a) = \mathbf{ta}_o^\diamond(\beta a)$.

DRM-3: We prove the contrapositive. Suppose $o \notin \text{alter}(\text{dom}(a), \alpha)$. We can assume that $o \in D$, since the values of objects in $\mathbf{0sets}$ do not change. Thus, we have $\alpha \models \text{dom}(a) \mapsto o$. Then we have $\text{ta}_o^\diamond(\alpha a) = \text{ta}_o^\diamond(\alpha)$ and hence $\text{contents}(o, \alpha a) = \text{contents}(o, \alpha)$.

DRM-4: Since the function **observe** is constant, we have $\text{observe}(u, \alpha \cdot a) \setminus \text{observe}(u, \alpha) = \emptyset$.

DRM-5: By definition, we have that

$$\begin{aligned} \text{observe}(u, \alpha) \cap \text{alter}(v, \alpha) &= \begin{cases} \{u\} & \text{if } \alpha \models v \mapsto u \\ \emptyset & \text{otherwise} \end{cases} \\ \text{observe}(u, \beta) \cap \text{alter}(v, \beta) &= \begin{cases} \{u\} & \text{if } \beta \models v \mapsto u \\ \emptyset & \text{otherwise} \end{cases} \end{aligned}$$

It is immediate that if $\alpha \models v \mapsto u$ and $\alpha \models v \mapsto u$ then $\text{observe}(u, \alpha) \cap \text{alter}(v, \alpha) = \{u\} = \text{observe}(u, \beta) \cap \text{alter}(v, \beta)$. (We remark that we do not need the antecedents $\alpha \approx_u \beta$ and $\alpha \approx_v \beta$ of DRM-5, so we have actually proved something stronger.)

DRM-6: We prove the contrapositive. Suppose $\alpha \models u \not\mapsto v$. Then clearly $\text{alter}(u, \alpha) \cap \text{observe}(v, \alpha) = \{w \in D \mid \alpha \models u \mapsto w\} \cap \{v, \text{oset}(v)\} = \emptyset$.

Next, suppose that M is \mathbf{ta}^\square -secure with respect to \mapsto . By Lemma 6 there exists a local policy $\mapsto' \leq \mapsto$ such that M is \mathbf{ta}^\square -secure with respect to \mapsto' . By Theorem 3.2, M is \mathbf{ta}^\diamond -secure with respect to \mapsto' . Hence, we may apply the above construction with respect to \mapsto' instead of \mapsto , to obtain an access control interpretation consistent with \mapsto' . We show that the stronger condition DRM-5' also holds in this case. Let $\alpha \approx_u \beta$ and $\alpha \approx_v \beta$. Then we have $\text{ta}_u^\diamond(\alpha) = \text{ta}_u^\diamond(\beta)$ and $\text{ta}_v^\diamond(\alpha) = \text{ta}_v^\diamond(\beta)$, i.e., $\alpha \sim_u^\diamond \beta$ and $\alpha \sim_v^\diamond \beta$. We may work with the relations \sim^\diamond in the definition of locality, so from locality of \mapsto' , it follows that $\alpha \models v \mapsto' u$ iff $\beta \models v \mapsto' u$. If $\alpha \models v \mapsto' u$ and $\beta \models v \mapsto' u$ then $\text{observe}(u, \alpha) \cap \text{alter}(v, \alpha) = \{u\} = \text{observe}(u, \beta) \cap \text{alter}(v, \beta)$. Alternately, if $\alpha \models v \not\mapsto' u$ and $\beta \models v \not\mapsto' u$ then $\text{observe}(u, \alpha) \cap \text{alter}(v, \alpha) = \emptyset = \text{observe}(u, \beta) \cap \text{alter}(v, \beta)$. That is, in either case, we have the desired result. \square

Theorem 5.2 can be understood as saying that every system that is secure with respect to a local policy could have been constructed using access control mechanisms satisfying the reference monitor conditions, in such a way that its security follows using Theorem 5.1. The phrase “could have been constructed” here is formally interpreted as allowing bisimilar systems as equivalent alternate implementations. Thus, this result states that the reference monitor conditions provide both a sound and complete methodology for the construction of systems that are secure with respect to local policies. This result, and the intuitive nature of the dynamic reference monitor conditions, provides further justification for our definitions.

6 Example: a capability system

In the operating systems literature, *capabilities* are a commonly used representation for access control. They provide a conceptual model that allows for dynamic security policies, by allowing an agent possessing a capability to transmit it to another agent, thereby granting the receiver the powers associated with the capability. A number of works on *Decentralized Information Flow Control* systems [21, 28, 35, 37] have proposed capability-based systems as a mechanism for implementing dynamic information flow policies. Many of the proposals in the literature contain covert channels. In this section, we develop a simplified model that is in the spirit of these proposals, but tailored to be secure with respect to an associated policy. (Our model was motivated by a consideration of Flume [21], but we have taken significant liberties in order to obtain a system that satisfies the desired result. We discuss the motivation for our changes to Flume in Section 6.4.)

6.1 States and Observations

We assume a set P of processes, each corresponding to a domain, so $D = P$. There is a fixed set \mathcal{N} of *basic tag names*, from which the set $\mathcal{T} = \mathcal{N} \cup \{n_p \mid n \in \mathcal{N}, p \in P\}$ of *tags* is generated. Here, each tag $t \in \mathcal{N}$ is a tag that exists at system startup, and n_p is a tag that is labelled with the process p that may create it. A security level is represented as a subset of \mathcal{T} . A *capability* is a value of the form either t^+ or t^- , where t is a tag. Intuitively, t^+ represents the capability to add tag t to a security level, and t^- represents the capability to remove tag t from a security level. We write *Capabilities* for the set of all capabilities.

Each process p has an associated set $Objects_p = Data_p \cup PolicyState_p$ of objects. The set $PolicyState_p$ represents process p 's component of the policy, and consists of the following special objects, each with an associated type (of possible values for the object)

- S_p of type $\mathcal{P}(\mathcal{T})$: the value of this object is a set of tags, and intuitively represents the types of information that process p may receive. The set S_p is called the *secrecy* set of process p , and represents its security level.
- O_p of type $\mathcal{P}(Capabilities)$ the value of this object is a set of capabilities. The set O_p is called the *capability set* of process p , and represents the set of powers that the process has to add and remove tags from its security level.

The set $Data_p$ is the set of data objects that are local to process p . We assume that the set $Data_p$ contains at least the following:

- m_p , an outgoing message buffer, used by p to construct a message prior to its transmission
- in_p , of type a sequence of messages. This is the input buffer through which p receives messages from other processes.

A *state* is a function s assigning a value of the appropriate type to each object in $\bigcup_{p \in P} \text{Objects}_p$. We write S for the set of states. A *candidate initial state* is a state s such that for all processes p , we have $O_p \subseteq \{n^x \mid n \in \mathcal{N}, x \in \{+, -\}\}$ and $S_p \subseteq \mathcal{N}$. That is, in an initial state, only basic tags exist; tags local to a process need to be created as the system runs. We allow the initial state of a system to be any candidate initial state: different choices of initial state correspond to different initial configurations of the security policy.

The observation functions of a capability system may be defined as any functions that satisfy the following constraint, for each process p and all states s and t :

Obs. If $s \upharpoonright \text{Objects}_p = t \upharpoonright \text{Objects}_p$ then $\text{obs}_p(s) = \text{obs}_p(t)$.

That is, observations depend, for each process, only on its own objects.

Given the intuition that the set S_p represents the types of information that process p may receive. Note that if S_p contains a tag t then p may have received information of type t . If t is not in S_q for some other process q , then there is a risk that if p transfers information to q then q will obtain information of type t , which is prohibited. This suggests that with each state s , we may naturally associate a static policy \mapsto , such that $p \mapsto q$ iff $S_p(s) \subseteq S_q(s)$. In what follows, it is convenient to write $s \models S_p \subseteq S_q$ when $s(S_p) \subseteq s(S_q)$.

6.2 Actions and Transitions

The following actions are possible in the system. We write actions in the format $p.a$, where $p \in P$ is a process, and define the domain of action $p.a$ to be $\text{dom}(p.a) = p$.

Data actions $p.a$

We assume that p has some set of data manipulating actions that act on its data objects. This set of actions can be instantiated in any way, subject to the constraint that these actions read only the local state of the process and write only the data objects of the process. More precisely, for any data action $p.a$, we assume the following, for all states s and t :

Data1. For all objects $x \notin \text{Data}_p$, $(s \cdot p.a)(x) = s(x)$.

Data2. If $s(x) = t(x)$ for all $x \in \text{Objects}_p$, then $(s \cdot p.a)(x) = (t \cdot p.a)(x)$ for all $x \in \text{Data}_p$.

Constraint **Data1** says that data actions of process p change only p 's data objects, and no others. Constraint **Data2** says that the way that such an action changes p 's data objects depends only on the value of p 's data and policy state objects. That is, a data action $p.a$ is deterministic and may read only objects in Objects_p and write only objects in Data_p .

$p.\text{add_cap}(c, n)$ where $n \in \mathcal{N}$ and $c \subseteq \{+, -\}$ with code

$$O_p := O_p \cup \{n_p^x \mid x \in c\}$$

Intuitively, this action creates a new tag n_p , if it does not already exist, and grants p the rights c over this tag, by adding the capabilities n_p^x for $x \in c$ to p 's capability set O_p . (If the tag already exists, the operation may be used to add capabilities to the existing ones.) Note that the fact that the tag created is indexed by the process p means that two distinct processes cannot create a tag with the same name.

$p.\text{drop_cap}(c)$ where $c \in \text{Capabilities}$, with code

$$O_p := O_p \setminus \{c\}$$

This action removes the capability c from p 's capability set, if present.

$p.\text{add_tag}(t)$ where $t \in \mathcal{T}$, with code

$$\text{If } t^+ \in O_p \text{ then } S_p := S_p \cup \{t\}$$

This action adds a tag to the secrecy set of the process, provided it has the capability to do so. A consequence of this is that some outgoing edges from p may be removed from the policy, and some ingoing edges may be added.

$p.\text{remove_tag}(t)$ where $t \in \mathcal{T}$, with code

$$\text{If } t^- \in O_p \text{ then } S_p := S_p \setminus \{t\}$$

This action removes a tag from the secrecy set of the process, provided it has the capability to do so, and consequently might add outgoing policy edges from p or insert incoming edges. Intuitively, by removing the tag, p enables declassification of information of type p ; whereas it may have received information of type t , the policy is no longer constrained to restrict p from communicating with processes that are not permitted to receive information of type t .

$p.\text{send_message_to}(q)$ where q is a process, with code

$$\text{If } S_p \subseteq S_q \text{ then } in_q := in_q \cdot m_p$$

This action transmits the message in the message object m_p to q by adding it to the end of q 's input buffer in_q , provided the policy permits transmission of information from p to q (i.e., $S_p \subseteq S_q$). The action has no effect on the policy, even if the message contains information about p 's tags and/or

capabilities. Intuitively, we assume that the operating system is responsible for maintaining the policy state, and mere knowledge of a capability name or value does not suffice to be able to exercise the capability. In order to transmit the ability to exercise capabilities, the following action needs to be used.

$p.\text{send_cap}(c, q)$ where $c \in \text{Capabilities}$ and q is a process, with code

If $S_p \subseteq S_q$ and $c \in O_p$ then $O_q := O_q \cup \{c\}$.

This action transmits the capability c to q , provided both that p is permitted to transmit information to q (i.e., $S_p \subseteq S_q$) and p actually possesses the capability (i.e., $c \in O_p$). The effect of this action may be to remove outgoing edges from q or add incoming edges to q .

6.3 Security of the capability model

We now establish the security of any capability system constructed as described above. Note that there are multiple such systems, since the description above leaves open the choice of processes, basic tag names, data objects, data actions, initial state and observation functions. It also remains to specify a dynamic security policy. Given a capability system M with actions A and processes P , and initial state s_0 , we say that the *associated policy* is the policy \rightsquigarrow^M defined by $\alpha \models p \rightsquigarrow^M q$ when $s_0 \cdot \alpha \models S_p \subseteq S_q$. The following states that a capability system is secure with respect to the the associated security policy.

Theorem 6.1 *Let M be a capability system. Then the associated policy \rightsquigarrow^M is local, and the system M is both \mathbf{ta}^\diamond -secure and \mathbf{ta}^\square -secure with respect to \rightsquigarrow^M .*

Proof: We show that M admits an access control interpretation consistent with \rightsquigarrow^M and apply Theorem 5.1.

For the set of objects of the access control interpretation, we take $O = \cup_{p \in P} \text{Objects}_p \cup \{\mathbf{oset}(p)\}$. (Since both the access control model and the capability model use a notion of *object*, and these are not quite the same when relating the two models, the reader needs to bear in mind that O refers to the objects in the access control model and *Objects* refers to objects in the capability model.) We define the access control functions for states s and processes p by $\mathbf{observe}(p, s) = \text{Objects}_p \cup \{\mathbf{oset}(p)\}$ and $\mathbf{alter}(p, s) = \text{Objects}_p \cup \{in_q, O_q \mid s \models S_p \subseteq S_q\}$. For each state s and object $o \in \cup_{p \in P} \text{Objects}_p$, we define $\mathbf{contents}(o, s) = s(o)$, and (as is necessary) for $o = \mathbf{oset}(p)$, we define $\mathbf{contents}(o, s) = \mathbf{observe}(s, p)$. We need to show that conditions DRM-1 to DRM-4, DRM-5' and DRM-6 are satisfied by these definitions.

Note that, by definition of \approx_p , and the definitions above, we have $s \approx_p t$ iff for all $o \in \text{Objects}_p$ we have $s(o) = t(o)$ (the contents of $\mathbf{oset}(p)$ are static).

Additionally, $o \in \mathbf{alter}(p, s) \cap \mathbf{observe}(q, s)$ iff $p = q$ and $o \in \text{Objects}_p$ or $p \neq q$ and $s \models S_p \subseteq S_q$ and $o \in \{in_q, O_q\}$.

DRM-1: This is immediate from the above characterization of \approx_p and the constraint **Obs**.

DRM-2: Suppose $o \in \text{alter}(\text{dom}(a), s) \cap \text{alter}(\text{dom}(a), t)$ and $s \approx_{\text{dom}(a)} t$ and $\text{contents}(o, s) = \text{contents}(o, t)$. We need to show that $\text{contents}(o, s \cdot a) = \text{contents}(o, t \cdot a)$. By definition, we have $o \in \text{alter}(\text{dom}(a), s) \cap \text{alter}(\text{dom}(a), t)$ just when either $o \in \text{Objects}_{\text{dom}(a)}$ or $o \in \{in_q, O_q\}$ for some process q such that $s \models S_{\text{dom}(a)} \subseteq S_q$. From $s \approx_{\text{dom}(a)} t$ we have that $s \upharpoonright \text{dom}(a) = t \upharpoonright \text{dom}(a)$, and from $\text{contents}(o, s) = \text{contents}(o, t)$ we have $s(o) = t(o)$.

We argue that $\text{contents}(o, s \cdot a) = \text{contents}(o, t \cdot a)$, i.e., $(s \cdot a)(o) = (t \cdot a)(o)$, for each possibility for the action a .

- If a is a data action $p.b$ (so that $p = \text{dom}(a)$), then if $o \notin \text{Data}_{\text{dom}(a)}$, we have by **Data1** that $(s \cdot a)(o) = s(o) = t(o) = (t \cdot a)(o)$. In the case $o \in \text{Data}_{\text{dom}(a)}$, we have $o \in \text{Objects}_{\text{dom}(a)}$, and $(s \cdot a)(o) = (t \cdot a)(o)$ follows by **Data2**.
- The code for each of the actions $p.\text{add_cap}(c, n)$, $p.\text{drop_cap}(c)$, $p.\text{add_tag}(t)$, $p.\text{remove_tag}(t)$, reads only objects in Objects_p and writes only objects in Objects_p . If a is one of these actions, then $p = \text{dom}(a)$ and the claim holds.
- The only object that can be changed by the code for the action $a = p.\text{send_message_to}(q)$ is in_q , so this is the only case of o that we need to consider. If $\text{dom}(a) = p = q$, then the condition $S_p \models S_q$ is a tautology, so the code performs the assignment $in_p := in_p \cdot m_p$ in both states s and t . If $\text{dom}(a) = p \neq q$, then from $o = in_q \in \text{alter}(\text{dom}(a), s) \cap \text{alter}(\text{dom}(a), t)$ we have both $s \models S_p \subseteq S_q$ and $t \models S_p \subseteq S_q$. Thus, again, the code takes the same branch in both cases, and since $m_p \in \text{Objects}_p$ we have $s(m_p) = t(m_p)$ and we have $(s \cdot a)(in_q) = s(in_q) \cdot s(m_p) = t(in_q) \cdot t(m_p) = (t \cdot a)(in_q)$, as required. The argument for the action $p.\text{send_cap}(c, q)$ is similar.

DRM-3: Suppose that $o \notin \text{alter}(\text{dom}(a), s)$. Since $\text{Objects}_{\text{dom}(a)} \subseteq \text{alter}(\text{dom}(a), s)$, we have $o \notin \text{Objects}_{\text{dom}(a)}$. It follows that data actions a and actions a of the forms $p.\text{add_cap}(c, n)$, $p.\text{drop_cap}(c)$, $p.\text{add_tag}(t)$, $p.\text{remove_tag}(t)$ do not change the value of o , since these actions change only the values of objects in $\text{Objects}_{\text{dom}(a)}$. In the case of $a = p.\text{send_message_to}(q)$ and $p.\text{send_cap}(c, q)$, we have from $o \notin \text{alter}(\text{dom}(a), s)$ that o is not one of the two possible objects in_q and O_q whose values these actions can change, so again, these actions do not change the value of o .

DRM-4: This is immediate from the fact that the value of **observe**(s, p) is a constant.

DRM-5': Suppose $s \approx_p t$ and $s \approx_q t$. We need to show that $\text{alter}(p, s) \cap \text{observe}(q, s) = \text{alter}(p, t) \cap \text{observe}(q, t)$. Now $\text{observe}(q, s)$ is a constant, and if $p = q$ then $\text{alter}(p, s) = \text{Objects}_p$ is also a constant, so we

need only consider the case $p \neq q$. In this case, $\text{alter}(p, s) \cap \text{observe}(q, s)$ is $\{in_q, O_q\}$ if $s \models S_p \subseteq S_q$ and is \emptyset otherwise. From $s \approx_p t$ and $s \approx_q t$ we get that $s(S_p) = t(S_p)$ and $s(S_q) = t(S_q)$, so $s \models S_p \subseteq S_q$ iff $t \models S_p \subseteq S_q$. Thus, $\text{alter}(p, s) \cap \text{observe}(q, s) = \text{alter}(p, t) \cap \text{observe}(q, t)$.

DRM-6: Note that for $p \neq q$, the sets $Objects_p$ and $Objects_q$ are disjoint. It follows that if $\text{alter}(p, s_0 \cdot \alpha) \cap \text{observe}(q, s_0 \cdot \alpha) \neq \emptyset$, then we have either $p = q$ or $s_0 \cdot \alpha \models S_p \subseteq S_q$. In either case, we have $s_0 \cdot \alpha \models S_p \subseteq S_q$, i.e., $\alpha \models p \mapsto^M q$. \square

It follows from Theorem 6.1 that a capability system is secure with respect to any policy that is no more restrictive than the associated policy.

Corollary 1 *Let M be a capability system. Then for any dynamic security policy \mapsto with $\mapsto^M \leq \mapsto$, the system M is both \mathbf{ta}^\diamond -secure and \mathbf{ta}^\square -secure with respect to \mapsto .*

Proof: Immediate from Theorem 6.1 using Proposition 2 and Proposition 3. \square

6.4 Comments on Flume

The capability system described above is somewhat motivated by Flume [21], but differs in a number of respects. We comment here on the differences.

First, Flume associates both a *secrecy set* $S_p \subseteq \mathcal{T}$ and an *integrity set* $I_p \subseteq \mathcal{T}$ to each agent p . These operate dually, so we have avoided inclusion of the integrity set just for simplicity. Our formulation could be extended to include it.

Next, Flume permits transmission of information from p to q when $S_p \setminus D_p \subseteq S_q \cup D_q$, where D_u is the set of tags for which u possesses both the capability to add and remove the tag. The intuition for this definition is that if p wishes to transmit information to q , but is constrained from doing so because S_p contains a tag t that are not in S_q , then this obstacle would be overcome if p is able to remove the tag from S_p or if q is able to add the tag to S_q . Assuming no other obstacles, the transmission could then proceed. Moreover, if, correspondingly, p is also able to add the tag or q is able to remove it, then after the transmission, the policy could be restored to its original state. Thus, the condition $S_p \setminus D_p \subseteq S_q \cup D_q$ captures the situations where this sequence of moves would allow transmission from p to q while leaving the policy state unchanged after the transmission is complete. The thinking then is that since such a sequence of moves cannot be prevented, the policy may as well allow the transmission directly.

We are not convinced that this is a good idea. One reason for p to have a tag in D_p but not in S_p might be that p wishes to protect itself against *inadvertent* transmission of information to q . Mechanisms such as *sudo* and the discretionary access controls on a user's own objects in Unix are explicitly designed to disable the effects of actions that the user nevertheless has the capacity to enable, and it seems useful to have a similar discretionary layer of control for information

flow policies. So we have not followed Flume in adopting this modification of the basic policy.

The Flume model includes a set \hat{O} of *global capabilities*, that are implicitly owned by all domains. That is, Flume works with the sets of owned capabilities $\overline{O}_p = O_p \cup \hat{O}$. When creating a tag, an agent is able to place the associated capabilities into \hat{O} . The motivation for this appears to be the following ([21] section II.D):

The most important policy in Flume is *export protection*, wherein untrustworthy processes can compute with secret data without the ability to reveal it. An export protection tag is a tag t such that $t^+ \in \hat{O}$ and $t^- \in \hat{O}$. For a process p to achieve such a result, it creates a new tag t and grants t^+ to the global set \hat{O} , while closely guarding t^- . To protect a file f , p creates the file with secrecy label $\{t\}$. If a process q wishes to read f , it must first add t to S_q , which it can do since $t^+ \in \hat{O} \subseteq \overline{O}_q$. Now, q can only send messages to other processes with t in their labels. It requires p 's authorization to remove t from its label or send data to the network.

We have not followed Flume in including the set of global capabilities \hat{O} . One reason is that we find the above argument unconvincing. Once t^+ is in \hat{O} , not just q , but any other agent r may add t to its secrecy set, and any protection afforded by tag t is then eliminated. The main reason that Krohn and Tromer have in mind that r might not be able to add t to its secrecy set is that t is randomly generated, and r does not know the appropriate value to use, making it a low probability event that r should add t unless it somehow receives t^+ in a chain of transmission from p . But in that case, we may as well rely upon the direct transmission of t^+ to the set O_r as the enabler for r 's adding t to its secrecy set. There does not appear to be any necessary role for \hat{O} in the above argument.

Another reason for not following Flume in the inclusion of the global set \hat{O} is that it causes a covert channel. An agent p may choose to add a capability t to O , and any other agent q can test whether the capability t is in O or not. This causes a flow of information from p to q . Krohn and Tromer are aware of this covert channel and seek to mitigate it by randomizing tag names. However this means that they require a much more complicated probabilistic version of noninterference semantics. It is not clear to us whether their semantics is well-formulated, since some subtle points appear to have been overlooked. Note that their systems have both probabilistic transitions (random generation of tag names) and nondeterministic transitions (agent's choice of their actions). This means that there is not a straightforward probability distribution on runs. How to handle the mix of probability and nondeterminism is an issue that has been studied in the literature, and involves a number of subtleties [18, 7]. It does not appear that Krohn and Tromer are aware of these issues, indeed, they do not even define the probability space that they use. It would seem that, at the very least, some non-trivial further details are needed to make their result fully

precise, which we will not pursue here.

Finally, Flume allows for dynamic creation of processes by means of a forking operation. It uses this to represent newly created objects as processes. Our policy model assumes a fixed set of domains, so this aspect of Flume cannot be directly represented. It would be interesting to develop an extension of our framework with the expressiveness to encompass creation of objects and agents.

7 Related Work

Related work using an automaton-based semantics similar to the one we have used in the paper has already been discussed in detail in Section 3.5. We focus here on work in the context of programming language security that has considered dynamic information flow policies. A recent survey of dynamic policies in the context of programming languages is [5]. They attempt to develop a classification based on various “facets” that are factored into the wide spectrum of definitions of security in this area. In general, there exist nontrivial gaps between the semantic frameworks used in this literature and our semantic model. Even *within* the programming language literature, formal relationships between the various semantics studied are rarely presented. We therefore give only some informal comparisons with our work.

One of the facets discussed in [5] is the distinction between “whitelisting flows” and “blacklisting flows”, which is similar to our distinction between a permissive and a prohibitive reading of policies. Interestingly, the whitelisting/permissive reading is claimed to be the norm in the programming language security literature.

However, there are also significant points of difference. Generally these works concern a programming language framework in which the secrets to be concealed are already encoded into the initial state, rather than our “interactive” model in which secrets are generated on the fly as the result of nondeterministic choices of action made by the agents. Frequently, the programming languages studied are deterministic, so a direct relationship to our setting is not immediately clear. There exists an approach by Clarke and Hunt [9] to handling interactivity in a programming language setting by including stream variables in the initial state to represent the sequence of future inputs that will be selected by an agent over the course of the run. However, this, in effect, assumes that the scheduler also is deterministic. Adding a variable for the schedule and allowing other agents to learn the complete schedule would be more permissive than a definition such as TA-security, since this restricts the information that agents may learn about the schedule. It therefore seems that establishing an exact correspondence would require some detailed work.

Another significant point of difference is that in the programming language literature, the policy and the system model (program) are often tightly intertwined, whereas we have a separate policy and system model, just as a specification and implementation are usually considered distinct. Other works that do allow for such a separation include Morgan [26] develops a refinement calculus

that allows for mixtures of information flow specifications and code, enabling an implementation to be derived from a specification via steps that mix the two, and Delft et. al. [33], who develop a type theory for dynamic information flow that operates on a program separately from the dynamic policy.

One of the ways that intransitivity is handled in the programming language setting also involves a dynamic aspect to policies: special constructs are introduced that allow temporary violations of a static policy, either through use of special basic actions or the establishment of a program region within which an alternate policy is enforced. For example, Mantel and Sands [23] work with a classical transitive security policy, but add downgrading commands that may violate this policy (subject to some alternate constraints.) They give a security definition that uses a specially defined bisimulation relation.

Broberg and Sands [3, 4] define *flow locks*, which enable variables to be typed with conditional annotations that constrain the security levels to which contents of these variables may flow. The values referred to in the conditions may be altered during the execution of the program, thereby making the policy dynamic. The authors give a complex bisimulation-based semantics to these policies [3], but in later work [4] prefer an alternate knowledge-based semantics, with a perfect recall interpretation of knowledge. It appears that their definitions are closer to those for a classical purge-based definition than to our intransitive policies. To capture a downgrader policy $H \rightsquigarrow D \rightsquigarrow L$, such as to allow information from H to reach L via D , it seems necessary to temporarily change the policy to effectively add an edge $H \rightsquigarrow L$ to the policy, allowing a direct flow from H to L . This is not in the spirit of intransitive noninterference, which intends to capture that all flows from H to L must pass through D .

Chong and Myers [8] propose to handle declassification violations of transitive policy using a notion of *noninterference until declassification* which applies a classical noninterference definition on runs just up to the point where a sequence of events has occurred on the basis of which the policy permits a declassification of secret information to the domain under consideration. Beyond this point, the definition does not constrain the behaviour of the system. Hicks et al [20] refine this idea to impose constraints on longer runs, declaring these secure if they consist of a sequence of policy-update free fragments, each of which satisfies the standard noninterference condition. This extension is also applied in Swamy et al [32], who develop a language with dynamic role-based information flow policies. This approach is also similar to the approach used by Almeida-Matos and Boudol [25], who work in a language-based setting where a transitive policy can temporarily be extended by addition of new edges (yielding another transitive policy). Their security condition states, using a bisimulation, that each transition in an execution is secure with respect to the transitive policy in place at the time the transition is taken. We note that these approaches generally do not handle covert channels resulting from the policy changes themselves, as illustrated in our Example 1.

All of the above approaches, like ours in this paper, work with perfect recall definitions of attacker knowledge. Askarov and Chong [1] point to an interesting issue with respect to dynamic security policies, which is that events may trans-

mit “old news”: information that would have been allowed to be transmitted to an attacker according to an earlier state of the security policy, but was actually not transmitted. The current state of the policy may prohibit the flow of this information. In this case, security definitions that are based on a perfect recall model of attacker knowledge, fail to identify such transmissions of “old news” as violations of security.

In response to this observation, Askarov and Chong propose to work with attackers that are weaker than the perfect recall attacker. They define a notion of knowledge for such attackers, and state definitions of security that require that any new knowledge gained by an attacker at security level (domain) u should only concern security levels v such that $v \rightarrow u$ according to the security policy at the time of the event. They identify some potential weaker attackers, but do not identify a definitive set of attackers; since the attacker is a parameter of their proposed security definition one could view the set of attackers against which a system should be secure as a matter of policy.

One of the attackers considered by Askarov and Chong is the perfect recall attacker, and here it is reasonable to attempt to compare their framework to ours. Unfortunately, making such a comparison formally is not completely straightforward because their semantic setting is somewhat different from ours: whereas in our framework there are multiple causal agents (the domains), who are free to choose any action at any time, in theirs there is a single causal agent (a deterministic program) that reads and writes information from different security domains. However, there is one respect where a clear difference of intent is apparent in their intuitions and definitions.

Consider a situation with domains $D = \{A, B, C, P\}$. Let Δ be the policy $\{u \rightarrow u \mid u \in D\} \cup \{P \rightarrow u \mid u \in D\}$. Intuitively, P here is the policy authority, and Δ states that all domains are permitted to know the state of the policy. In the initial state of the system, take the policy to be $\Delta \cup \{A \rightarrow B\}$. Any actions performed by A are, intuitively, permitted by this policy to be recorded in the state of B , but should remain unknown to C . Suppose that after some actions by A , the policy agent changes the policy to $\Delta \cup \{B \rightarrow C\}$. Now, B is permitted to communicate with C . According to our intuitions and formal definitions, B may now send to C the information it has about the actions that A performed before the policy change. Such flows from A to C via B are typical of the flows that intransitive policies are intended to permit [17].

However, according to Askarov and Chong’s definitions (both for the perfect recall attacker and weaker attackers), if an action now copies this information from B ’s state and it is observed by C , this is a violation of security, because in a transition where it makes the observation, C learns new information about A , whereas the current state of the policy prohibits it from doing so. Thus, although Askarov and Chong permit intransitive policies, we would argue that their definitions do not handle dynamic intransitivity in a way that fits our intuitions: their semantics can best be characterized as a dynamic version of the classical purge-based definition of security, rather than a dynamic version of a semantics for intransitive noninterference. The details of Askarov and Chong’s semantics are refined in [33], but similar remarks apply to this work.

8 Conclusion

We have developed two generalizations for intransitive dynamic policies of the notion of TA-security that has recently been argued to provide a better semantics for static policies. One is based on a permissive reading of policy edges, the other on a prohibitive reading of policy edges. Under the natural condition of locality, the two readings are equivalent. We have also provided proof theory for these notions in the form of unwinding conditions and a sound and complete policy enforcement method using access control. Finally, we have demonstrated the applicability of the theory by applying it to establish security for a natural capability model.

Several interesting directions remain open for further investigation. One concerns the circularity in the justification of \mathbf{ta}^\square -security. While \mathbf{ta}^\square -security is, as we have shown, self-consistent, there may be further natural constraints that also lead to self-consistent notions of security. For example, one candidate is to require for information transfer that the sender know that there is an edge from it to the receiver, in place of the assumption we made in \mathbf{ta}^\square -security that the sender and receiver must jointly know that there is such an edge. Whether such a constraint should apply is arguably a matter of policy. This suggests that there may be a case for a policy language that more explicitly provides a way to express conditions on agent knowledge. It would also be of interest to develop algorithms for verifying whether a system satisfies a policy. Both models and dynamic policies as we have defined them are highly expressive, so this would require some restrictions: finite state models and policies that are representable using finite state automata would be one interesting starting point: it remains to be shown that this case is decidable.

We have shown that while \mathbf{ta}^\diamond -security is more expressive than \mathbf{ta}^\square -security, the difference between these notions applies only in the case of non-local policies. Both are associated with natural intuitions, but we leave open the question of whether the extra expressiveness of \mathbf{ta}^\diamond -security is necessary in practice, i.e., whether there are any interesting systems requiring non-local policies where \mathbf{ta}^\diamond -security, rather than \mathbf{ta}^\square -security, is the appropriate notion of security. The main realistic example we have presented, the capability system of Section 6 involves a local policy, so it does not resolve this question. The notion of \mathbf{ta}^\diamond -security allows information flows relating to policy settings that are not explicitly represented as edges in the policy itself. Conceivably, a richer policy format may allow these implicit flows to be more explicitly represented.

It would also be interesting to formally relate our work to work in the programming language setting, and to combine the insights from Askarov and Chong’s work on imperfect recall attackers with the approach we have developed in the present paper.

References

- [1] A. Askarov and S. Chong. Learning is change in knowledge: Knowledge-based security for dynamic policies. In S. Chong, editor, *CSF*, pages 308–322. IEEE, 2012.
- [2] C. Boettcher, R. DeLong, J. Rushby, and W. Sifre. The MILS component integration approach to secure information sharing. In *Proc. 27th IEEE/AIAA Digital Avionics Systems Conference*, pages 1.C.2–1–1.C.2–14, Oct. 2008.
- [3] N. Broberg and D. Sands. Flow locks: Towards a core calculus for dynamic flow policies. In P. Sestoft, editor, *ESOP*, volume 3924 of *Lecture Notes in Computer Science*, pages 180–196. Springer, 2006.
- [4] N. Broberg and D. Sands. Flow-sensitive semantics for dynamic information flow policies. In *PLAS*, pages 101–112, 2009.
- [5] N. Broberg, B. van Delft, and D. Sands. The anatomy and facets of dynamic policies. In *IEEE 28th Computer Security Foundations Symposium, CSF 2015, Verona, Italy, 13-17 July, 2015*, pages 122–136, 2015.
- [6] D. Bytschkow, J. Quilbeuf, G. Igna, and H. Ruess. Distributed MILS architectural approach for secure smart grids. In *Smart Grid Security - Second International Workshop, SmartGridSec 2014, Munich, Germany, February 26, 2014, Revised Selected Papers*, pages 16–29, 2014.
- [7] L. Cheung. *Reconciling Nondeterministic and Probabilistic Choices*. PhD thesis, IPA, Radboud University, Nijmegen, 2006.
- [8] S. Chong and A. C. Myers. Security policies for downgrading. In *Proceedings of the 11th ACM Conference on Computer and Communications Security, CCS 2004, Washington, DC, USA, October 25-29, 2004*, pages 198–209, 2004.
- [9] D. Clark and S. Hunt. Non-interference for deterministic interactive programs. In *Formal Aspects in Security and Trust, 5th International Workshop, FAST 2008, Malaga, Spain, October 9-10, 2008, Revised Selected Papers*, pages 50–66, 2008.
- [10] E. S. Cohen. Information transmission in computational systems. In *SOSP*, pages 133–139, 1977.
- [11] S. Eggert, H. Schnoor, and T. Wilke. Noninterference with local policies. In K. Chatterjee and J. Sgall, editors, *MFCS*, volume 8087 of *Lecture Notes in Computer Science*, pages 337–348. Springer, 2013.
- [12] S. Eggert, R. van der Meyden, H. Schnoor, and T. Wilke. Complexity and unwinding for intransitive noninterference. *CoRR*, abs/1308.1204, 2013.

- [13] K. Engelhardt, R. van der Meyden, and C. Zhang. Intransitive noninterference in nondeterministic systems. In *ACM Conference on Computer and Communications Security*, pages 869–880, 2012.
- [14] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT-Press, 1995.
- [15] J. A. Goguen and J. Meseguer. Security policies and security models. In *IEEE Symposium on Security and Privacy*, pages 11–20, 1982.
- [16] J. A. Goguen and J. Meseguer. Unwinding and inference control. In *IEEE Symposium on Security and Privacy*, 1984.
- [17] J. T. Haigh and W. D. Young. Extending the noninterference version of MLS for SAT. *IEEE Transactions on Software Engineering*, SE-13(2):141–150, Feb 1987.
- [18] J. He, K. Seidel, and A. McIver. Probabilistic models for the guarded command language. *Sci. Comput. Program.*, 28(2-3):171–192, 1997.
- [19] C. L. Heitmeyer, M. Archer, E. I. Leonard, and J. D. McLean. Formal specification and verification of data separation in a separation kernel for an embedded system. In *Proc. of the 13th ACM Conf. on Computer and Communications Security, CCS*, pages 346–355, 2006.
- [20] M. Hicks, S. Tse, B. Hicks, and S. Zdancewic. Dynamic updating of information-flow policies. In *Proc. of Foundations of Computer Security Workshop (FCS)*, 2005.
- [21] M. Krohn and E. Tromer. Noninterference for a practical DIFC-based operating system. In *IEEE Symposium on Security and Privacy*, pages 61–76, 2009.
- [22] R. Leslie. Dynamic intransitive noninterference. In *Proc. IEEE International Symposium on Secure Software Engineering*, 2006.
- [23] H. Mantel and D. Sands. Controlled declassification based on intransitive noninterference. In *Proc. Asian Symp. on Programming Languages and Systems*, volume 3302 of *LNCS*, pages 129–145. Springer-Verlag, Nov. 2004.
- [24] W. Martin, P. White, F. Taylor, and A. Goldberg. Formal construction of the mathematically analyzed separation kernel. In I. C. S. Press, editor, *Proc. 15th IEEE Conf. on Automated Software Engineering*, 2000.
- [25] A. A. Matos and G. Boudol. On declassification and the non-disclosure policy. *Journal of Computer Security*, 17(5):549–597, 2009.
- [26] C. Morgan. The shadow knows: Refinement and security in sequential programs. *Sci. Comput. Program.*, 74(8):629–653, 2009.

- [27] T. Murray, D. Matichuk, M. Brassil, P. Gammie, and G. Klein. Noninterference for operating system kernels. In C. Hawblitzel and D. Miller, editors, *Certified Programs and Proofs*, volume 7679 of *Lecture Notes in Computer Science*, pages 126–142. Springer Berlin Heidelberg, 2012.
- [28] A. C. Myers and B. Liskov. A decentralized model for information flow control. In *Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles*, SOSP '97, pages 129–142, New York, NY, USA, 1997. ACM.
- [29] A. W. Roscoe and M. H. Goldsmith. What is intransitive noninterference? In *IEEE Computer Security Foundations Workshop*, pages 228–238, 1999.
- [30] J. Rushby. Noninterference, transitivity, and channel-control security policies. Technical Report CSL-92-02, SRI International, 1992.
- [31] G. Schellhorn, W. Reif, A. Schairer, P. A. Karger, V. Austel, and D. C. Toll. Verified formal security models for multiapplicative smart cards. *Journal of Computer Security*, 10(4):339–368, 2002.
- [32] N. Swamy, M. Hicks, S. Tse, and S. Zdancewic. Managing policy updates in security-typed languages. In *19th IEEE Computer Security Foundations Workshop, (CSFW-19 2006), 5-7 July 2006, Venice, Italy*, pages 202–216, 2006.
- [33] B. van Delft, S. Hunt, and D. Sands. Very static enforcement of dynamic policies. In *Proc. Principles of Security and Trust - 4th Int. Conf., POST*, pages 32–52, 2015.
- [34] R. van der Meyden. What, indeed, is intransitive noninterference?1. *Journal of Computer Security*, 23(2):197–228, 2015. An earlier version of this work appeared in ESORICS'07.
- [35] S. Vandebugart, P. Efstathopoulos, E. Kohler, M. Krohn, C. Frey, D. Ziegler, F. Kaashoek, R. Morris, and D. Mazières. Labels and event processes in the Asbestos operating system. *ACM Trans. Comput. Syst.*, 25(4), Dec. 2007.
- [36] W. Vanfleet, R. Beckworth, B. Calloni, J. Luke, C. Taylor, and G. Uchenick. MILS:architecture for high assurance embedded computing. *Crosstalk: The Journal of Defence Engineering*, 2005.
- [37] N. Zeldovich, S. Boyd-Wickizer, E. Kohler, and D. Mazières. Making information flow explicit in HiStar. In *7th Symposium on Operating Systems Design and Implementation (OSDI '06), November 6-8, Seattle, WA, USA*, pages 263–278, 2006.